Characterizing Laser Signal Injection and its Impact on the Security of Cyber-Physical Systems

by

Benjamin Andrew Cyr

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Computer Science and Engineering)
in The University of Michigan
2023

Doctoral Committee:

Professor Kevin Fu, Northeastern University, Co-Chair
Professor Mingyan Liu, Co-Chair
Professor Atul Prakash
Associate Professor Alanson Sample
Professor Herbert Winful

Benjamin Andrew Cyr

bencyr@umich.edu

ORCID iD: 0000-0003-1986-3554

To Sierra:

Another Step on our Shared Adventure

ACKNOWLEDGEMENTS

It has been a long, difficult, but ultimately rewarding journey, and I want to thank everyone who helped me to reach my goals.

Family and Friends: First, I want to thank my family and friends for all of the support over my academic career. I want to thank my parents, brother, grandparents, aunts, uncles, cousins, and my wife's parents for always supporting my efforts and providing me with new opportunities to further my goals. You have always been there when I needed support, and I cannot express how much I appreciate being able to ask for help with anything. I want to also thank my best friend, Mark Gallagher, who was a brother to me and always available to chat about research or life in general. I want to thank all of those who helped me keep my sanity outside the lab: Charlie, Kellen, Crowe, Allison, Anna, Riley, Nick, Riley, Angel, Julia, and Aditya. And of course, I want to thank my amazing wife, Sierra Cyr, who was my companion even in the most stressful of times.

SPQR Lab: Next, I want to thank the current and past members of the SPQR Lab: my advisor Kevin Fu, who set me up with so many opportunities; Sara Rampazzi and Connor Bolton, both my friends and mentors in the graduate program and life in general; Yan Long, who always gave insightful advice and discussion as we worked together on our discoveries; all of the undergraduates, master's students, and visiting students who assisted in our projects and were passionate to make discoveries together.

Collaborators: I want to thank all the collaborators who helped develop the re-

search ideas presented in this thesis: Daniel Genkin, Z. Morley Mao, Wayne Burleson, Srinivas Tadigadapa, Yulong Cao, Vedant Sumaria, Kohei Yamashita, Chaowei Xiao, Yimeng Zhou, Won Park, Qi Alfred Chen. I especially want to thank Takeshi Sugawara, who advised me on many of the ideas within these works.

University of Michigan: I also want to thank the many individuals at the University of Michigan who are always willing to have discussions and give advice: Karl Krushelnick, John Nees, Karl Grosh, Mark Brehob, Alanson Sample, Atul Prakash, Herbert Winful, and Mingyan Liu.

Teachers: Finally I want to thank the teachers, professors, and advisors that inspired me to continue down this path not just by sharing knowledge, but through their passion for discovery and learning: Yvette Kromann, Edward Thomas, Victor Nelson, Robert Dean, Dean Hendrix, J.M. Wersinger, Yin Sun, and Ujjwal Guin.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

LIST OF ABBREVIATIONS

LFI Laser Fault Injection

LSI Laser Signal Injection

TD Thermal Diffusion

TE Thermoelastic Effects

PV Photovoltaic Effects

MEMS Micro-electro-mechanical Systems

LiDAR Light Detection and Ranging

APD Avalanche Photodiode

AV Autonomous Vehicle

VCS Voice-Controllable System

ASIC Application-Specific Integrated Circuit

IR Infrared

ADACS Attitude Determination and Control Subsystem

EPS Electrical Power Subsystem

MPPT Maximum Power Point Tracking

ABSTRACT

Lasers can be used to inject adversarial-controlled signals into sensors used in cyber-physical systems. This capability is often unexpected, use physical mechanisms that were never considered, and exploits the blind trust in sensors. These laser signal injection attacks can allow adversarial influence or even control over a system's perception of the environment, leading to potentially harmful situations.

The contributions of this work are the characterizations of laser signal injection in three cyber-physical contexts: LiDAR sensors used in autonomous vehicles, MEMS microphones in voice-controllable systems, and the sensors used in space systems. These characterizations include an in-depth investigation of attacker capabilities, the development of models to describe the vulnerability, and a description of the consequences on the relevant cyber-physical systems.

The characterization of laser signal injection into LiDAR sensors used in autonomous vehicles builds upon previous research to define the capabilities of an attacker to influence LiDAR data. This work shows that false laser returns can be spoofed at a large number of angles and formed into specific shapes. These shapes can be used to trick object detection algorithms to register false objects, causing autonomous vehicles to make dangerous control decisions. This is followed by recommendations for future defenses by adjusting the object detection algorithms or making modifications to the LiDAR sensors to prevent laser signal injection.

The characterization of laser signal injection into MEMS microphones used in voice-controllable systems is the investigation of a previously unknown vulnerability. This work shows that MEMS microphones are unexpectedly vulnerable to laser signal

injection attacks due to a combination of photoelectric and photoacoustic phenomena. Models of these effects are presented, as well as a comparison of the relative contributions of these effects on the output of the microphone given different laser injection characteristics. Beyond this, there is also an investigation of the vulnerabilities of voice-controllable systems that rely on microphone data to perform autonomous actions, leading to new threat models. Defense recommendations for protecting the microphones at the system level or developing new light-resistant microphones are presented.

The characterization of laser signal injection into sensors used in space systems is a preliminary investigation of future threat models. This work includes a preliminary attack surface analysis of sensors that are potentially vulnerable to laser signal injection. It goes over the current capabilities of a potential adversary as well as capabilities on which to focus future research to determine future vulnerabilities. An example case study on light-sensitive sun sensors is presented to show an example attack on space sensors. Beyond this, recommendations are made to set up test benches for future vulnerability research on satellites, as well as recommendations for future defenses to detect laser single injection.

These characterizations are performed not just to present the attacks, but to fully understand the mechanisms that lead to vulnerabilities within these cyber-physical systems. This understanding is necessary to design future systems in a way that will be resistant to all forms of laser signal injection, allowing sensors and cyber-physical systems to be safer, more trustworthy, and more secure.

CHAPTER I

Introduction

For all of human history, humans have relied upon the senses of sight, hearing, touch, smell, and taste in order to perceive the world, understand its current state, and take actions that affect that state. Now technology has developed to the point where we can provide our tools with our own senses in the form of sensors. Some of these sensors capture the same information as our senses, but even more capture information well outside of our capabilities. These sensors have enabled the development of new tools that no longer rely on human senses but can extend our perception, make their own judgments about the current state of the environment, and even make computational decisions to interact with the physical world. These cyber-physical systems extend our capabilities and improve our lives, and sensors are essential to their operation.

Fundamentally, sensors rely on the capture of environmental energy, whether that be mechanical energy, thermal energy, electromagnetic energy, or any other form of energy. Electromagnetic energy in the form of light is particularly important due to the way it propagates through space. It is ubiquitous, easy to measure and capture, doesn't require a medium, and can provide a significant amount of information about the environment. But these same features that make light an ideal mechanism to capture information about the environment also make it an ideal mechanism to change

a system's perception of the environment.

Light can be used by a malicious adversary to control a system's perception of the environment through its sensors. With control of perception, the adversary ultimately influences the system's control decisions, potentially pushing an automated system to a state that can be harmful. When this capability is paired with the directionality and energy density of a laser, it opens up a wide range of potential targets vulnerable to laser signal injection (LSI) attacks.

1.1 Motivation

There is a semantic gap between the way system designers expect a sensor to work, and how the sensor actually measures the environment. Systems are designed around the idea that the sensors will only perform in a certain way: a LiDAR sensor will only give distance measurements to the surrounding environment, a microphone will only measure audible signals, and a sun sensor will only give you a vector in the direction of the sun. But researchers have shown that sensors can be affected in ways that were never expected: sensors can "over sense" the environment in a way that introduces erroneous data, and systems that blindly trust these sensors are insecure because of it [2, 3]. While the reliability field continues to develop ways to prevent environmental noise from causing system failures, the prevention of intentional and malicious signal injection is still an open research problem. These signal injection attacks on sensors need to be understood to develop defenses for future devices.

There are major consequences to signal injection on sensors. Every autonomous system has sensors that allow it to perform its functions. When an attacker controls the sensors, they control a system's perception of reality, and ultimately some control over the system itself. In autonomous vehicles, an adversary that can inject false signals into the sensors used in perception can jeopardize the safety of people and property around the vehicle. This opens up the possibility for attackers to gain

access to users' private data or finances with devices such as smart speakers. In satellites and space systems, an adversary that can control sensor data can cause mission failures, damage expensive equipment, or pose threats to national defense. As these automated systems continue to become more ubiquitous, defending against signal injection attacks is going to become even more important.

1.2 Thesis Statement and Summary

Sensors utilized in current cyber-physical systems are vulnerable to laser signal injection attacks in unexpected ways; therefore, a strong understanding of these current vulnerabilities is necessary to ensure that future defenses can be expected to keep systems safe and secure. This is the fundamental idea motivating this thesis and research. The purpose of this thesis is to investigate and characterize laser signal injection attacks on cyber-physical systems. This involves exploring the capabilities of potential adversaries of the present and the future, developing physical models that describe potential attacks, measuring the potential consequences on cyber-physical systems, and recommending defenses based on these discoveries. Ultimately the research focus is to better understand the threats and vulnerabilities that exist or are likely to exist in the near future so that future systems can be developed to be resistant to laser signal injection.

1.3 Contributions

In this thesis, I describe my research contributions toward characterizing laser signal injection into the sensors that are used in cyber-physical systems. These contributions are:

- A characterization of laser signal injection into LiDAR sensors used in autonomous vehicles (Chapter III). This includes a description of attacker ca-

pabilities, a method to generate coherent spoofed data, an investigation into the system consequences, and recommendations for future defenses and open problems. This research has resulted in a publication at the 2019 ACM Conference on Computer and Communications Security [4].

- A characterization of laser signal injection into MEMS microphones used in voice-controllable systems such as smart speakers and smartphones (Chapter IV). An investigation of an attack on voice-controllable systems research was presented at the 2020 Usenix Security Conference in the form of *Light Commands* [5]. A further investigation of the causality of LSI in MEMS microphones led to a publication at the IEEE SENSORS 2021 conference [6]. Beyond these works, the resulting model of LSI in MEMS microphones, recommendations for defenses, and the remaining areas of research are presented.

- A characterization of laser signal injection into sensors used in space systems (Chapter V). This includes a preliminary description of the potential attacks with LSI on satellite sensors, the current capabilities of a potential adversary, and an example case study on light-sensitive sun sensors. This research resulted in a publication in the 2023 Workshop on Security of Space and Satellite Systems [7].

CHAPTER II

An Overview of Laser Signal Injection

Laser Signal Injection (LSI) is the use of light energy in the form of lasers to add signals within a system in order to test, disrupt, or control the output of the system. It is commonly used in reliability testing and fault detection, as it provides a very precise way to inject a highly controlled signal into a localized part of the system. More recently, however, there has been an increased focus on the consequences laser signal injection will have on the security of vulnerable systems. The purpose of my research is to characterize some of the risks that laser signal injection poses to these systems.

To do this, it is important to begin by explaining previous work and the important considerations that are common to laser signal injection. In this chapter, I will discuss the principles of signal injection, the physical phenomena that have been explored to inject energy into a system via laser, and the previous works that have studied laser signal injection for uses in reliability and security.

2.1 Signal Injection Attacks

A signal injection attack is the addition of an adversarial signal to an important signal within a system in an effort to disrupt or control the perception of the signal. This adversarial noise is generated by a transfer of energy from an adversarial sys-

tem to the victim system through a variety of physical mechanisms: direct electrical injection, intentional electromagnetic interference (IEMI), acoustic vibrations, lights and lasers, and potentially more. From the perspective of the system, the adversarial signal is considered "noise": an unwanted addition to the signal that changes its perception of a true parameter. But unlike random environmental noise that occurs naturally in any environment, adversarial noise can be intelligently crafted to exploit vulnerabilities within the system. That is what makes signal injection so dangerous.

Signal injection attacks often exploit vulnerabilities within an important input to any system: sensors. Sensors are already designed to be sensitive to environmental signals and are therefore often sensitive to other forms of adversarial energy transfer. Most devices today rely upon a fundamental assumption: sensors provide a signal about the state of the system of the environment within a certain expectation of noise. This blind trust in sensors leads to all kinds of problems, both in reliability and security.

As described in Yan et. al. [8], the sensor can be modeled as a chain of signal processing transfer functions that converts a state x into a measured value y. Each transfer function represents a component of the sensor: transducers, amplifiers, filters, analog-to-digital converters, etc. The combination of each component in this chain is described by:

$$y = f(x; N) \tag{2.1}$$

where f is the total sensing transfer function and N is the set of environmental noise signals that get added to the measurement at each sub-component. The noise set N can be represented by:

$$N = \{n_1, n_2, n_3, \ldots n_m \ : n_i \sim P_i(n)\} \tag{2.2}$$

where each member n_i is a random sample from some distribution of environmental noise $P_i(n)$. Each n_i is added to the measurement as it propagates through the sensor components. The output of the sensor y is then the combination of the state x and the transformations of all the noise sources in N. The system then uses a function g to determine an approximation of the current state x from the measurement y.

$$g(y; E[N]) = x + \varepsilon \tag{2.3}$$

The function g is dependent on the expected values of environmental noise $E[N]$, as the goal of the function is to minimize the noise error ε. In well-designed systems, the environmental noise N is greatly attenuated, and the knowledge about the noise is strong enough that the function g produces a strong approximation of the current state x.

$$g(y; E[N]) \approx x \tag{2.4}$$

In the case of a signal injection attack, an adversarial noise signal A is injected into certain components of the sensor, summing with the environmental noise. This means the measured signal y' can be modeled as:

$$y' = f(x; N + A) \tag{2.5}$$

When the system attempts to approximate the state x based on the measurement y':

$$g(y'; E[N]) = x + \alpha + \varepsilon \approx x + \alpha \tag{2.6}$$

there is an additional term α that changes the perceived state due to the adversarial signal. While g can be designed to remove the expected environmental noise and minimize ε, the distribution of A can be anything within the capabilities of the

attacker. When there is no prior knowledge of the distribution of A, there is no way to remove the α term, and the system perceives the current state as $x + \alpha$. This additional term allows an adversary to control the system's perceived state, and potentially control the output of the system.

The purpose of this work is to give a stronger knowledge about the capabilities to inject the adversarial signal A. With better knowledge and characterization of the adversarial capabilities, future systems can be designed to reduce or detect the presence of the α term, ultimately defending against signal injection attacks.

2.2 Transferring Energy via Lasers

As described above, signal injection relies upon the transfer of energy from the adversarial system to the victim system. Light is one of the most efficient ways to transfer energy long distances without any infrastructure or medium. Direct electrical injection requires an electrical connection and potential access to the victim system. Acoustic injection requires a mechanical connection and often relies on the vibrations through the air, which poorly couples to solid components. The long wavelengths of the electromagnetic waves used in IEMI mean that the energy diffuses quickly, greatly reducing the energy absorbed by the target at even short ranges. The short wavelengths of optical light allow for a transfer of energy at high densities and precision.

In particular, lasers provide a precise way to inject this energy into the target system. In this section, I will discuss how lasers function and the important parameters that affect the amount of energy that lasers transfer to the target system. These properties of lasers and their effects are important in understanding the capabilities of LSI attacks.

2.2.1 Laser Overview

A laser is a device that uses stimulated emission and optically resonant cavities to generate light with properties that are ideal for signal injection. The acronym LASER stands for "Light Amplification by Stimulated Emission of Radiation". Stimulated emission is the key phenomenon that is being used, which was first hypothesized by Albert Einstein in 1917 [9], where photons are used to stimulate the emission of more photons out of a charged medium. Stimulated emission provides a feedback mechanism to amplify the light signal and greatly increase its power density.

An example laser system is shown in figure. Lasers consist of an optical resonator cavity filled with some kind of gain medium. External energy is used to excite the atoms within the gain medium via a number of different mechanisms. Once light of a certain wavelength enters the cavity, it resonates, stimulating the emission of more photons at a similar wavelength and amplifying the optical signal. By providing a transmission mechanism out of the resonator cavity, a laser beam is generated to be potentially used in LSI.

In general, lasers have four properties that make them useful as a mechanism for signal injection [10]:

- **Directionality**: Light from conventional light sources tends to diverge quickly due to the unstructured light emission and the size of the light-generating element. While optics can be used to improve the directionality, the limiting factor of the divergence is the geometry of the source and the optics. The limiting factor for lasers, however, is diffraction, where the wave nature of light causes it to spread out as it propagates. The divergences due to diffraction are much smaller, allowing the light to be tightly focused and travel long distances. This will be explored more in Section 2.2.3.

- **Spectral Purity**: Unlike conventional light sources, lasers are capable of pro-

ducing very small spectral bandwidths consisting primarily of a single color. This is due to the optical resonator design, where only a small set of optical wavelengths resonate and achieve amplification. This is useful in signal injection, as many injection mechanisms are sensitive to the wavelength of the incoming light.

- **High Power**: Because of their strong directionality and amplification properties, lasers can be used to transfer high power to a relatively small area at range. Many of the signal injection mechanisms are proportional to the optical irradiance (i.e. optical power per unit area), making lasers an ideal candidate for signal injection.

- **Controllability**: Lasers can output optical power in a way that is highly controllable. Some continuous-wave lasers produce an output that is linearly proportional to an easily controlled parameter such as current. Pulsed lasers can be tuned to produce very short pulses with very precise timing. While there are always limitations on the control for every laser system, these capabilities allow an attacker to craft a precise signal that can be used to inject data into a target system.

These four properties make lasers an ideal tool for long-distance signal injection.

2.2.2 Semiconductor Lasers and Modulation

While laser signal injection is possible with all lasers, this work uses semiconductor lasers to investigate signal injection, as they are low-cost and easily accessible. Semiconductor lasers use doped semiconductors as a gain medium. By organizing them into a p-n junction as a diode, an electrical current can be used to generate charge carriers within the depletion region of the diode. These charge carriers will spontaneously emit light over time as they recombine, but can also be used in stimu-

lated emission. With the edition of special optical structures such as Bragg gratings, an optical resonator can be formed within the depletion region, amplifying light of particular wavelengths and generating a laser beam that propagates out of the diode.

Semiconductor laser diodes are easily controlled using a current source. Each pair of charge carriers displaced by the current has a chance to recombine and emit a photon within the p-n junction. Once enough current is present, enough photons are emitted to cause stimulated emission, greatly increasing the output optical power of the diode. The current at which stimulated emission occurs is called the threshold current (I_{th}:A). Below the threshold current, the output power is roughly equal to zero. Once the current is greater than the threshold, the number of photons, and therefore the optical power output (P_0:W) is linearly proportional to the input current flowing across the junction diode (I_{in}:A). This relationship is modeled by:

$$P_0 \propto \begin{cases} I_{in}, & I_{in} > I_{th} \\ 0, & I_{in} \leq I_{th} \end{cases} \tag{2.7}$$

and some examples of this relationship are shown in Figure 2.1. Since the irradiance exiting the attacker's laser system is proportional to the optical power output, a voltage-controlled current source provides a convenient way to linearly transform any voltage signal into an optical irradiance signal, provided the diode current is greater than the threshold current.

2.2.3 Beam Propagation

One of the most unique properties of the use of lasers in signal injection is its directionality. The power density of lasers can be considerable even at long distances, and they can be focused into tight beams. At relatively close ranges, this allows an attacker to inject energy to precise locations on vulnerable targets, having much better control over the signal being injected. At long ranges, power can be transmitted

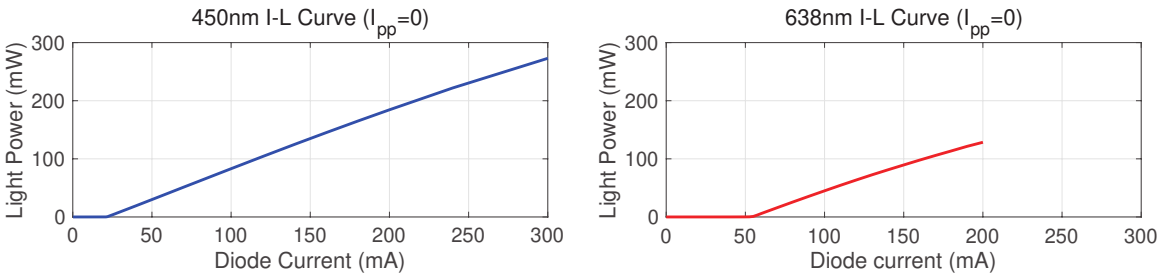

Figure 2.1: Current vs output optical power for two semiconductor lasers: a 450 nm Osram PLT5 450B (left) and a 638 nm Thorlabs L638P150 (right). The optical output power is linearly proportional to the input current once the current is greater than some threshold value.

from the attacking system to the vulnerable system much more efficiently than nearly any other signal injection mechanism. It is important to understand how laser beams propagate through space to understand the important parameters that determine the limits of an attacker's capabilities. Much of the equations and models are taken from Thyagarajan and Ghatak [10], Self [11], and Andrews et. al. [12]

Ideal lasers are diffraction-limited when they are collimated and have the minimum divergence possible. Diffraction is the process by which light spreads and interferes with itself as it propagates through space. It is often used as an illustration of Heisenberg's Uncertainty Principle [13], as the divergence angle (θ) of the beam is approximately proportional to the inverse of the beam waist, where the beam waist (w_0: m) is the radius of the beam at its most focused point. In other words, the more is known about the position of the photons, the less is known about the momentum and the beam has a greater divergence.

The divergence angle of the laser due to diffraction at long distances is:

$$\theta \approx \frac{\lambda}{\pi w_0} \tag{2.8}$$

where λ is the wavelength of the electromagnetic wave in meters and w_0 is the beam waist in meters. This model shows how the beam divergence is greater with smaller

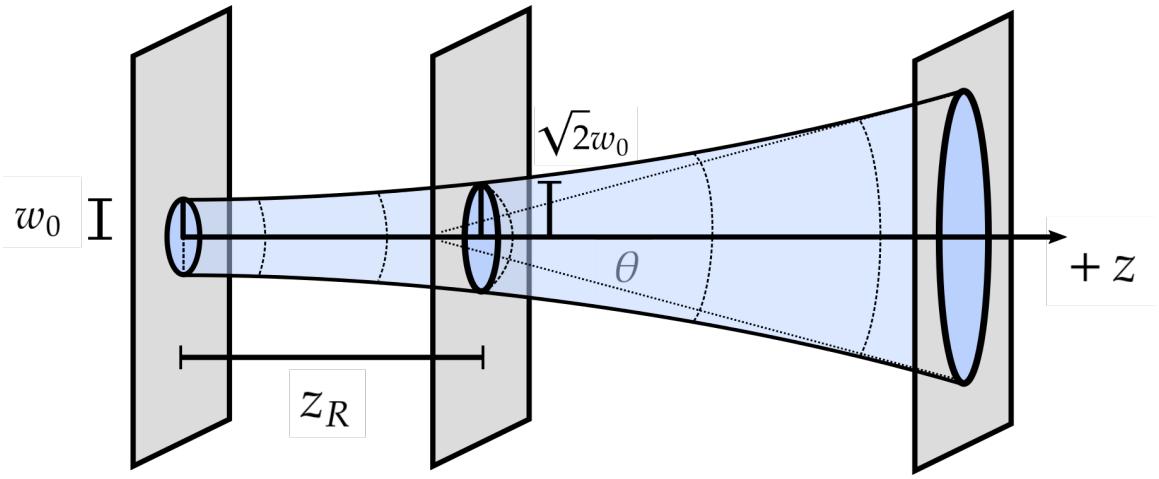

Figure 2.2: The propagation of a Gaussian laser beam through space. The beam stays fairly collimated, with a radius roughly equal to the beam waist (w_0) until the Rayleigh Range (z_R), where the diffraction angle (θ) increases significantly

beams and longer wavelengths of light. Therefore, to use a laser signal for long-range injection, large optics, and small light wavelengths will be most effective.

To further define how a laser beam propagates through space, a Gaussian electromagnetic beam model can be used. The radius (w: m) of a circular Gaussian beam is defined as the radial distance between the center of the beam and where the intensity of the light is at $1/e^2$ of the center. If the output of a laser system is perfectly collimated, the initial beam radius will be at the minimum, making it the beam waist (w_0: m). The beam radius will change as the wave propagates a distance z in the following way:

$$w(z) = w_0\sqrt{1 + \left(\frac{z}{z_R}\right)^2} \tag{2.9}$$

$$z_R = \frac{\pi w_0^2}{\lambda} \tag{2.10}$$

where z_R is the Rayleigh range, a special distance that dictates the transition between a "near-field" case where the irradiance stays fairly constant, and the "far-field" case

where the beam diverges based on the diffraction angle described in Eq. 2.8. The Rayleigh range is an important parameter that indicates when the beam radius is a factor of $\sqrt{2}$ larger than the beam waist.

We can then model the irradiance of the beam using the beam waist and the Rayleigh range:

$$I(r,z) = \frac{I_0}{1+(z/z_R)^2} exp\left(\frac{-2r^2/w_0^2}{1+(z/z_R)^2}\right) \tag{2.11}$$

$$I_0 = \frac{2P_0}{\pi w_0^2} \tag{2.12}$$

where r and z are the radial and distance coordinates, and I_0 is the initial irradiance of the center of the beam at the output of the optical system. The initial irradiance is defined from the initial area of the beam and the total power of the optical beam.

Notice that at distances much less than the Rayleigh range, the irradiance is not significantly affected, meaning the beam stays roughly collimated. At the Rayleigh range exactly, the irradiance is exactly half of the initial irradiance. For distances much greater than the Rayleigh range, the z/z_R term dominates the denominator, and the beam diverges via its diffraction divergence angle. The irradiance in the far field can be approximated by:

$$I(0,z) \approx I_0\left(\frac{z_R^2}{z^2}\right) = I_0\left(\frac{\pi^2 w_0^4}{z^2\lambda^2}\right) for z \gg z_R \tag{2.13}$$

One can clearly see that the irradiance follows an inverse-square law and quickly drops in intensity once reaching this far-field condition.

For the purpose of LSI, the goal of the attacker is to maximize the irradiance striking the target device. The most effective way to accomplish this is to maximize the Rayleigh range of the optical system. This can be done by using shorter wave-

lengths of light and optics such as beam expanders to make the beam radius as large as possible on the output of the injection system.

2.2.4 Material Reflectance, Transmittance, and Absorption

While the effectiveness of a LSI attack is proportional to the irradiance of incoming light, it is important to understand how the optical properties of the target materials influence potential LSI threats. These optical properties will determine how much of the energy is actually transferred from the adversarial system to the target system.

As an electromagnetic wave in the form of a laser interacts with a boundary of a material, a certain amount of the incoming power is reflected off of the boundary, and the rest transmits into the material. The portion of the light that is reflected or transmitted is highly dependent on the wavelength of the incoming light and the optical properties of the materials at the interface. The reflectance (R_λ) of a material boundary is defined as its effectiveness at reflecting incoming radiant. It is a value defined from zero to one, where zero is a purely transmitting boundary, and one is a purely reflecting boundary. The reflectance of an interface can be modeled by the well-known Fresnel Equations with the assumption that the incoming light is normal to the boundary [14]:

$$R_\lambda = \left(\frac{n_{\lambda 1} - n_{\lambda 2}}{n_{\lambda 1} + n_{\lambda 2}} \right)^2 \tag{2.14}$$

where $n_{\lambda 1}$ is the refractive index of one material at the boundary, and $n_{\lambda 2}$ is the refractive index of the second material. Note that the refractive index for every material is dependent on wavelength λ of the incoming light, meaning that the reflectance is also dependent on wavelength. The remaining proportion is defined as the transmittance and transmits past the boundary into the material.

$$T_\lambda = 1 - R_\lambda \tag{2.15}$$

Once the transmitted light enters the material, it begins being absorbed by the particles within the material. The incoming power increases the kinetic energy of the particles, resulting in increased temperature and the potential to increase the energy levels of a particular particle. The irradiance as the light travels through the material can be modeled via the Bouguer-Beer-Lambert Law [15]:

$$I(z) = I_0 T_\lambda e^{-\beta_\lambda z} \tag{2.16}$$

where I_0 is the irradiance of the incident light at the surface, and β_λ is the absorption coefficient of the material. Note that the absorption coefficient is different for each material, and it is also dependent on wavelength. For most opaque materials and structures, a large absorption coefficient will result in most of the energy being absorbed near the surface, and the irradiance will drop to zero before transmitting through the structure.

In the case of transparent structures and thin films, however, the absorption coefficient is small enough so that the light will not be completely absorbed. Once it reaches a second boundary, some energy will be reflected back towards the first surface and the rest will transmit out of the material. As the electromagnetic waves travel back and forth within the thin film, it interferes with itself, leading to complex optical properties depending not just on the absorption coefficient, but also the thickness of the structure. The film itself acts as a resonant cavity for particular electromagnetic wavelengths and an antiresonant cavity for others. When multiple of these thin film structures are stacked, it greatly increases the complexity even more, making it very difficult to determine the optical properties analytically.

All of these factors influence the feasibility of LSI, as the optical properties of the materials in the target system determine the effectiveness of shielding and the efficiency of energy absorption. The wavelength of the laser signal is very important in LSI due to the wavelength dependence of the optical properties of the various

materials. Beyond this, certain thin structures such as ones present in MEMS devices will have very complex optical properties, which can make it difficult to predict how LSI will affect them.

2.3 Mechanisms of Laser Signal Injection

Numerous mechanisms exist to inject optical laser energy into a vulnerable system. The purpose of this section is to describe these different mechanisms and show how they all relate to the incoming optical irradiance, which is controlled by the adversary in LSI. The different mechanisms can be roughly organized into phenomena that rely on the direct conversion of optical energy into electrical energy, phenomena that rely on the conversion of optical energy to thermal energy, and phenomena that rely on the momentum transfer of photons. These mechanisms are explored below as mechanisms potentially useful for LSI, but it is not an exhaustive list.

2.3.1 Photoelectric Phenomena

Photoelectric phenomena are phenomena that rely on the excitation of electrons in response to incoming photons. This can be exploited in various ways to change various electrical properties or even generate mechanical motion. The phenomena are discussed below:

2.3.1.1 The Photoelectric Effect

The photoelectric effect is the process by which high-energy photons are absorbed by electrons within a material, imparting enough kinetic energy so that the electrons break their atomic bonds and eject from the material. As the electrons are stripped from the material, it becomes positively charged and can disrupt the electrical conditions within the material and any attached circuit in various ways.

First described by Heinrich Hertz in 1887 [16], the photoelectric effect was foundational to the development of quantum theory. Albert Einstein hypothesized in 1905 [17] that the photoelectric effect was caused by light "quanta" as discussed by Max Planck in 1901 [18]. In it, a quantum of light was defined to have an energy:

$$E = h\nu = hc/\lambda \tag{2.17}$$

where h is Planck's constant ($\approx 6.626 \times 10^{-34} J \cdot s$), ν, λ, and c are the frequency (Hz), wavelength (m), and speed (m/s) of the light wave. If an incoming photon has enough energy to overcome the "work function" of the material then the electron would be ejected. The work function is then the energy required to ionize the atom of the material, and any remaining energy would be converted to the kinetic energy of the electron.

While this work was instrumental to quantum theory and led to a Nobel prize in 1923, the photoelectric effect has limited applications within the LSI attacks investigated in this work. This is because the majority of common electrical materials have work functions that require spectral photon energies in the ultraviolet range ($\lambda < 400nm$). While a potential area of future research, the focus of the laser signal injection attacks in this work will be in the visible or infrared spectrum, where photons do not have enough energy to generate the photoelectric effect.

2.3.1.2 The Photovoltaic Effect

The photovoltaic effect is the generation of voltage and current within a material exposed to incoming optical energy. Similar to the photoelectric effect, the photovoltaic effect relies upon the absorption of photons by electrons within a material. Unlike the photoelectric effect, however, the electrons do not eject from the material but instead transition from a low-energy state to a high-energy state. Depending on

the chemical properties, geometry, and attached circuitry, this transition leads to a photovoltaic signal that can be used for optical sensing in photodiodes, power generation in solar cells, or an injection mechanism in LSI. Many of the details about photovoltaic effects in this section are described by Honsberg and Bowden [19].

The majority of work in modern photovoltaics focuses on semiconductors such as silicon. Within these semiconductors, incoming light energy can excite an electron from a valence band (which is tightly bound to the nucleus of an atom) to a conduction band (where the electron can move freely around the material). The amount of energy required to excite the electron is called the "band gap". The excited electron and the hole in the valence band are referred to as "charge carriers" or "electron-hole pairs". In pure semiconductors, these charge carriers move around randomly through the material for some lifetime before eventually recombining via a number of mechanisms and dumping the energy as heat or light. This is shown in Figure .

When dopants are added to semiconductors, different photovoltaic signals can be generated. By doping semiconductors in a particular way to form a P-N junction, charge carriers generated by incoming light behave differently than in pure semiconductors. Instead of random motions that lead to recombination, the inherent electric field of the p-n junction causes the charge carriers to drift towards opposite ends of the junction. In isolation, this generates a voltage as opposite charges collect on opposite ends of the junction. When connected to a circuit, the junction behaves as a current source as the photo-generated charges travel to the rest of the circuit.

In the photovoltaic effect, the amount of photocurrent (I_Φ: A) that is generated is proportional to the incoming photon flux (Φ : $photons/s/m^2$) and the quantum efficiency (η_λ):

$$I_\Phi \propto \eta_\lambda \Phi \tag{2.18}$$

The photon flux is the number of photons striking the p-n junction per unit area per second and can be calculated from the irradiance ($I : W/m^2$) and the energy ($E : J$)

of the individual photons within the incoming light.

$$\Phi = I/E = \frac{I\lambda}{hc} \tag{2.19}$$

The relationship between photon flux and photocurrent arises from the fact that each photon can normally only excite a single charge-carrier pair. This means that for similar irradiances, longer wavelengths of light will generate more photocurrent than shorter wavelengths, as the photon flux is directly proportional to wavelength.

Quantum efficiency is the parameter that determines how effectively the p-n junction absorbs photons of a particular wavelength, as the junction geometry and material absorption are significant factors that determine the probability that a generated charge carrier will recombine before drifting out of the junction. As an idealized approximation, the quantum efficiency can be modeled as unity for photon energies greater than the band gap of the semiconductor and zero for energies less than the band gap. For silicon, with a band gap of approximately $1.11eV$, the cutoff wavelength is approximately $1100nm$.

$$\eta_\lambda \approx \begin{cases} 1, & \lambda < 1100nm \\ 0, & \lambda \geq 1100nm \end{cases} \tag{2.20}$$

From the previous equations, it can be seen that the photocurrent generated will be proportional to the wavelength of the incoming light up to a cutoff wavelength of light, where no more charge carriers will be generated.

The photovoltaic effect is the mechanism that is most commonly exploited in LSI. Nearly all complex electrical systems depend upon semiconductors and p-n junctions to form diodes and transistors. Every p-n junction is potentially sensitive to some wavelengths of incoming light via the photovoltaic effect, generating photocurrents that can short logic circuits, flip bits in memory, or inject additional electrical signals

into analog systems. The research works investigating the photovoltaic effect as an injection mechanism are explored further in Section 2.4.

2.3.1.3 Photoconductivity

Outside of p-n junctions, the generation of charge carriers via incoming light will also change the conductivity of a bulk semiconductor material. First discovered in selenium by Willoughby Smith in 1873 [20], photoconductivity has been a key principle in many photodetectors. The change in conductance (ΔG: Ω^{-1}) of a semiconductor will be proportional to the excess electron and hole concentrations (Δn, Δp: $carriers/m^3$) generated by the incoming light [21]:

$$\Delta G \propto \mu_n \Delta n + \mu_p \Delta p \tag{2.21}$$

where μ_n and μ_p are the electron and hole drift mobilities. Since the excess electron and hole concentrations are proportional to the photons striking the semiconductor, the change in conductance of the semiconductor will therefore be proportional to the incoming photon flux (Φ: $photons/s/m^2$).

$$\Delta G \propto \eta_\gamma \Phi \tag{2.22}$$

Photoconductivity is an important factor in sensing systems that rely on resistive sensing of semiconductor components, as the resistance is the inverse of the conductance ($R = 1/G$). This can allow direct injection on light-dependent resistors (LDRs), or indirect injection on a component such as semiconductor piezoresistors. Incoming light will dynamically change the resistance, which provides a mechanism for LSI. Photoconductivity has only been explored as an LSI mechanism in a few cases, but can potentially be an important factor to consider to prevent future LSI attacks.

2.3.1.4 Plasmaelastic Deformation and Bending

While the previous mechanisms result in changes to electrical properties, plasmaelastic effects generate changes to the mechanical properties of a material. As charge carriers are generated within a semiconductor material in response to light, the shift in charge distribution causes a volume change within the crystal structure of the material. In the case of silicon, this actually results in a contraction of the material as the atoms within the crystal structure pull tighter together. This volume change can form stresses within the material in a similar way to thermal expansion, which can lead to mechanical motion and vibration. Ultimately, this means plasmaelastic effects can be considered a photoacoustic phenomenon, where modulated light is used to generate acoustic vibrations.

The "electronic volume effect" in semiconductors was first described by Gauster and Habing [22] in 1967, and later referred to as the "concentration-deformation mechanism" [23], the "electronic strain mechanism" [24], and finally the "plasmaelastic mechanism" [25]. The plasmaelastic mechanism is modeled in a similar way to thermal expansion, where a change in material length (ΔL: m) is dependent on the initial length (L_0), a linear coefficient of plasmaelastic deformation (d_n: $m^3/carrier$), and the concentration of excess minority carriers (Δn: $carriers/m^3$):

$$\Delta L = d_n L_0 \Delta n \tag{2.23}$$

The linear coefficient of deformation for silicon is negative ($\approx -9\times10^{-31} m^3/carrier$ [25]), but other semiconductors have a positive coefficient of deformation that results in an expansion.

As the volume of the semiconductor structure changes in response to the incoming light, it can also generate stress gradients that cause bending within semiconductor structures. This is especially a concern when thin insulating films are deposited on

semiconducting substrates. The substrate that is undergoing plasmaelastic deformation puts stress on the unaffected insulating film, causing the structure to bend. Depending on the structure, the displacement due to bending can be many orders of magnitude greater due to the displacement due to the volume change directly. Because of this, most works investigating plasmaelastic effects focus on the bending of mechanical structures such as plates [26] and cantilevers [27, 28].

Just like in previous photo-generated carrier effects, the concentration of excess carriers is proportional to the number of incoming photons. This indicates that any mechanical displacement due to plasmaelastic deformation or bending (w_{PE}) will also be proportional to the incoming photon flux (Φ: $photons/s/m^2$).

$$w_{PE} \propto \eta_\lambda \Phi \tag{2.24}$$

The plasmaelastic mechanism is a lesser-known way for optical energy to be transferred to a victim system. Because it is only applicable to semiconductors, it will only be a potential mechanism for LSI in cases where a mechanical semiconductor structure is within line-of-sight of an adversarial laser signal. While this is rare in conventional devices, the rise of micro-electro-mechanical systems (MEMS) has made it much more common for small, semiconductor structures to be influenced by incoming light. In general, however, this effect will be competing with thermoelastic effects, which tend to be stronger.

2.3.2 Photothermal Phenomena

Photothermal phenomena are phenomena that rely on the generation of heat as optical energy is absorbed by a material. In general, an increase in temperature (ΔT: K) of any material due to an absorbed laser beam will be proportional to the irradiance of the beam (I:W/m^2).

$$\Delta T \propto I \tag{2.25}$$

This conversion of optical energy into thermal energy can have multiple effects on a system by changing electrical properties or inducing mechanical motion. This can potentially be used to inject a signal into a sensitive system. Due to the relatively slow nature of heat, however, it is often difficult to generate precise signals except at low frequencies unless the target device has a very low heat capacity. Despite, this, there are some cases where this thermal phenomenon can be exploited. The different thermal effects are listed here to discuss their potential use in LSI.

2.3.2.1 Thermoresistive Effects

Many electrical properties within circuits are highly dependent on the temperature of the components. This is because the temperature, which describes the average kinetic energy of the particles, greatly affects how electrons and holes drift through the material in the presence of an electric field. The primary way the circuits are affected is by changes in resistance.

Resistors are sensitive to temperature. In most conductors such as copper, the increase in temperature increases the number of collisions experienced by electrons in the wire and increases the overall resistivity of the material. For semiconductors, however, the conductivity and therefore resistivity is dependent on the concentration of charge carriers within a semiconductor. While in photoconductivity (Section 2.3.1.3) these charge carriers are generated by incoming photons, higher temperatures will also generate more charge carriers within the semiconductor, lowering the resistivity at increasing temperatures. This mechanism has been studied since its discovery by Michael Faraday in 1833 [29] and has led to the development of using resistors as a temperature sensing mechanism.

Many resistors can be used as thermal sensors called “thermistors”. While com-

plex models of thermistors are often used [30], many thermistors often follow a linear model for the range of temperatures of interest. The change in resistance (ΔR: Ω) can be modeled with:

$$\Delta R = \alpha_R \Delta T \tag{2.26}$$

where ΔT is the change in temperature in Kelvin and α_R is the temperature coefficient of resistance. The temperature coefficient can be both positive (PTC) or negative (NTC) which is an important factor in designs that use thermistors.

While temperature-sensitive thermistors present an obvious location for the precise heating capabilities of LSI, it is less known that many resistors that are not designed as thermistors can be exploited in LSI. This is especially true for semiconductor resistors such as piezoresistors, which will have some inherent thermally-sensitive properties. Systems that rely on sensing these temperature-dependent resistors may be vulnerable to LSI.

2.3.2.2 Thermoelectric Effects

Besides thermoresistive effects, there are two other thermoelectric phenomena that can potentially be exploited:

First is the Seebeck effect, where a temperature difference between two ends of material will generate a potential and a current between the two ends. This effect is the result of the increased kinetic energy of charge carriers within the hot region of the metals diffusing into the cold end. This is often modeled with a linear model and a Seebeck coefficient (S: V/K):

$$\Delta V = -S \Delta T \tag{2.27}$$

where ΔV is the change in voltage and ΔT is the change in temperature. The Seebeck

effect is often used in sensitive thermocouple designs to measure temperature.

Second is the pyroelectric effect where sudden changes in temperature generate a transient electrical potential within a material. Highly related to the piezoelectric effect, the pyroelectric effect is generated by the sudden change in the distribution of charges within an asymmetric crystal. This effect is often used in passive infrared (PIR) sensors, as even the radiated heat given off by living creatures can generate a measurable signal within pyroelectric devices.

Similar to thermistors, these two effects can be exploited directly by LSI by heating up the sensitive component, whether that be a thermocouple, PIR sensor, or any other sensor using these effects. There may also be cases where piezoelectric materials such as the ones used in certain sensor MEMS designs may also exhibit pyroelectric sensitivity, which may be exploited in LSI.

2.3.2.3 Thermoelastic Expansion and Bending

As a solid material is heated by a laser, the average kinetic energy of the particles within the material increases. With the increased kinetic energy, the particles push against their bonds, causing an expansion within the material structure. Depending on the absorbing structure, these expansions can generate displacements, stresses, and elastic waves that can affect the functioning of a system, providing a way to exploit it via LSI.

The change in length of a material (ΔL: m) under a change in temperature (ΔT: K) is often described by a linear model [31]:

$$\Delta L = \alpha_T L_0 \Delta T \tag{2.28}$$

where L_0 is the initial length, and α_T is the linear coefficient of thermal expansion. This is modeled in a similar way to plasmaelastic effects in Section 2.3.1.4, though

nearly all solids will have a positive coefficient of thermal expansion.

If the material is constrained in any way, the change in length due to thermal effects will generate stresses within mechanical structures. This occurs both when the structure is heated, but also when it is cooled. While in general, the process of thermal expansion is considered a slow process, the stresses induced by thermal expansion can be considerable. Thermal stresses are often a concern in MEMS designs, as the manufacturing process often requires high temperatures, and residual thermal stresses as the device cools can lead to poor performance or even damage the MEMS structure [32].

Similar to plasmaelastic effects, thermal expansion can generate stress gradients that cause bending. This occurs when there is some asymmetry within the thickness of a structure. An asymmetry can be caused either by a temperature gradient through the thickness, or by mismatches in coefficients of thermal expansion of materials within a multilayer stack. The asymmetry results in a stress gradient that causes the structure to bend and generate a potentially greater displacement than with pure linear expansion. Both thermoelastic expansion and bending combine linearly, and the total displacement of the heated structure (w_{TE}) will also be proportional to the change in temperature (ΔT).

$$w_{TE} \propto \Delta T \tag{2.29}$$

Signal generation via laser heating and thermal expansion has been greatly explored within the context of photoacoustics. Photoacoustics is the study of the conversion of optical energy into mechanical vibrations. While there are many mechanisms to achieve this, photoacoustic signals generated by thermoelastic expansion of solids were first modeled by R. M. White in 1962 [33]. In this work, White showed how the sudden thermal expansion of a solid will generate an elastic wave that propagates through the material. McDonald and Wetsel developed this theory further in

1978 [34], showing how this elastic wave causes displacements and acoustic signals in the air. The signals generated by this mechanism are often fairly small but can provide a way to induce vibrations within a structure. Rousset et. al. [35] continued the exploration in 1983 by exploring how the asymmetric heating generated by laser absorption in opaque plates often leads to thermoelastic bending. This signal can often be significantly stronger than the pure expansion signal. Finally, photoacoustic signal generation in the case of thin films and mismatches in thermal properties was explored by Tódorović et al. in 2013 [25].

Thermoelastic expansion and bending from LSI is primarily a concern in systems that are sensitive to mechanical stress and strains leading to displacements. This is often a concern in MEMS devices that measure mechanical motion to a very fine degree. The structures in MEMS are also very small, which means that even a small amount of incoming energy is enough to increase the temperature to generate measurable expansions.

2.3.2.4 Thermal Diffusion

More than the thermal expansion of solids, the indirect heating and expansion of a gas can also be used as a mechanism for LSI. Within the context of photoacoustics, this process is often referred to as "thermal diffusion", as it relies upon the relatively slow diffusion of heat from a thermal source to a surrounding gas. Because the heat transfer is slow, a small volume of gas near the heated surface will be at a significantly higher temperature than the rest of the gas. If the temperature change is sudden, the small volume expands adiabatically, pressing against the rest of the gas and generating an acoustic pressure wave. This pressure wave can interact with various structures within the system to cause mechanical motion.

Thermal diffusion is considered the first photoacoustic phenomenon to be discovered [36] while Alexander Graham Bell and Sumner Tainter were developing a "photo-

phone" in 1880 to send audio information with light. While their initial investigations focused on using the photoconductive properties of selenium, they discovered that an audible acoustic signal would be generated by focused light even without electrical signals [37]. A working model of this photoacoustic mechanism was not developed until the "thermal-piston" model of Rosencwaig and Gersho in 1976 [38], which was later referred to as the "thermal diffusion" mechanism [35]. This model was incorporated into every following investigation into photoacoustic effects.

The fundamental principle of thermal diffusion is the sudden temperature increase of a thin layer of gas near a heated surface generates an adiabatic expansion within the gas. This boundary layer of gas acts as a piston, pressing against the rest of the gas, which is modeled as an ideal gas. This adiabatic expansion combined with the ideal gas law produces the following relationship between the change in pressure (ΔP_{TD}: Pa) and the temperature change of the thin layer of gas (ΔT_g: K) [38]:

$$\Delta P_{TD} = \frac{\gamma P_0 V_g}{V_0 T_0} \Delta T_g \tag{2.30}$$

where V_g P_0, V_0, and T_0 are the ambient temperature, volume, and temperature of the surrounding gas, and V_g is the volume of the heated gas layer. Note that because any displacement due to thermal diffusion (w_{TD}: m) will be roughly proportional to the pressure amplitude, and the change in temperature of the boundary gas will be proportional to the change in temperature of the solid surface (ΔT: K), the displacement will be proportional to the temperature change of the heated target.

$$w_{TD} \propto \Delta T \tag{2.31}$$

For small optical powers and low frequencies, thermal diffusion is most often the strongest photoacoustic effect, meaning it is one of the most effective photothermal signal generation mechanisms. Because of this, thermal diffusion is one of the

strongest LSI mechanisms when a vulnerable mechanical component is surrounded by gas. In most cases, however, the mechanical coupling between the acoustic pressure wave and the vibration of the solid structure is too weak to be effective. The exception is with very small mechanical structures such as the ones present in MEMS devices.

2.3.2.5 Laser Ablation

Extremely high-intensity lasers can thermally vaporize and ionize a material into plasma, ejecting particles from the surface. Due to the conservation of momentum, these ejected particles will impart pressure upon the material [39], which can generate acoustic vibrations. This process is referred to as laser ablation. It is most effective with short laser pulses, where most of the energy remains at the surface of the target material. This effect was first described in a presentation by Brech and Cross in 1962 [40], and later described in detail by John Ready in 1971 [41]. This is a highly nonlinear process, but some analytical model relates the laser ablation pressure (P_{LA}: Pa) to the irradiance (I: W/m^2)as [42]:

$$P_{LA} = P_0 I^{\alpha} \tag{2.32}$$

where P_0 is some initial pressure dependent on the properties of the laser and material, and α is a scaling exponent dependent on the properties of the plasma. The exponent α has normal values between 2/3 and 3/4.

Laser ablation is only applicable in cases with extremely high irradiances or very short pulses, as it requires a significant amount of energy in a small area. If this condition is met, however, it can generate very strong pressures and efficiently transfer optical energy into acoustic energy. Due to the large energies required, however, complex setups and future research will be required to determine its application to LSI.

2.3.3 Radiation Pressure

The final phenomenon that can be used to transfer energy into a vulnerable system is radiation pressure. Radiation pressure is the transfer of momentum of photons within a beam of light onto a material that reflects or absorbs the beam of light. Albert Einstein derived the expression relating the momentum of photons to their wavelength ($p = h/\lambda$) [43]. Due to the conservation of momentum, whenever the photons are absorbed by or reflected off of a material, the momentum of the photons is transferred to the material. This momentum transfer manifests as a tiny pressure (P_{RP}: Pa) applied to the material. This can be modeled with the following expression:

$$P_{RP} = (1 + R_\lambda)\frac{I}{c} \tag{2.33}$$

where R_λ is the reflectance, I is the incoming irradiance, and c is the speed of light. Note that the reflectance increases the effect, as more momentum is transferred by reflected photons than absorbed photons. Because the speed of light is so fast, this pressure is often extremely small, and most explorations of the effect require very sensitive equipment.

In most cases, radiation pressure will be unable to provide a significant signal to affect a target system. In fact, the primary cases where radiation pressure is a concern are in space, where the lack of resistive forces means that the tiny pressure will add up over a long time. From the context of LSI, this will only be applicable in cases of extremely small targets or extremely high irradiances.

2.4 Related Work

Many previous works served as an inspiration and foundation for my research into LSI. These works have explored how signals and faults can be injected into a system, showing the consequences on the reliability security of many different systems.

2.4.1 Laser Fault Injection

Laser Fault Injection (LFI) is the use of lasers to generate faults in computing devices via photoelectric effects. By injecting a pulse of laser energy at a very precise location on an exposed semiconductor chip, an adversary can flip bits to compromise data and logic flow, extract security keys, or dump device memories. Most LFI attacks use the photovoltaic effect to induce a photocurrent within the p-n junctions that form the transistors that make up the semiconductor chips. The photocurrent can be used to generate voltages or drain current in precise parts of the chip, giving substantial control to an attacker. But some works have also relied on photothermal effects such as thermoresistivity and the Seebeck effect [44], indicating that there are a number of ways lasers can be used to interact with systems.

Habing [45] showed that the photovoltaic effect can be used to test semiconductor devices by generating electrical transients with infrared light. Building off of this, Skorobogatov and Anderson [46] used lasers to flip individual bits of a memory cell, enabling laser fault attacks on smartcards and microcontrollers. This type of memory fault injection was used by Agoyan et al. [47] to show the vulnerability of cryptographic systems to lasers. Some work such as [48] have provided methods to detect LFI to indicate the attack to the chip logic. Finally, I helped with some results in *Redshift* [49], where lasers were used to influence signal propagation within PUF timing circuits to determine a secret key. Other LFI attacks and defenses are summarized by [50] and [51].

2.4.2 Cyber-Physical Security of Sensor-reliant Systems

My research on LSI stems from previous research into the security of sensors used in cyber-physical systems. This field of research focuses on how various physical mechanisms can be used to inject signals into systems reliant on sensors, determining vulnerabilities, and suggesting defenses. Many different injection mechanisms have

been explored, including IEMI, acoustic signals, and light. This has led to the discovery of many potential vulnerabilities and the development of new countermeasures to defend against signal injection. A formal review of cyber-physical sensor attacks and defenses was made by Yan et al. [52].

2.4.2.1 IEMI Signal Injection

A foundational work in this field was *Ghost Talk* [53], where IEMI was used to inject signals to sensors on implanted defibrillators, pacemakers, and microphones from a distance. This work inspired significant effort into finding other IEMI vulnerabilities within other systems reliant on sensors. In a work by Kasmi and Esteves [54], microphones with long wires were exploited to inject voice commands into a smartphone. Works by Chauhan [55] and Yan et. al. [56] demonstrated that some automotive radar devices are vulnerable to IEMI spoofing and jamming signals. The magnetic sensors used in automotive anti-locking braking systems were also shown to be vulnerable to magnetic spoofing [57]. Selveraj et. al. [58] showed analog sensors reliant on analog-to-digital converters (ADCs) were vulnerable to IEMI attacks. Beyond this, they also showed how sufficiently high IEMI can also inject signals into digital communication systems for both sensors and actuators. In *Tap 'N Ghost* [59], *Invisible Finger* [60], and *GhostTouch* [61] all used IEMI to induce fake touches in capacitive touch screens in tablets and smartphones. In *Trick or Heat* [62], IEMI was used to inject analog signals on thermocouples and other temperature-sensing systems by abusing a rectification effect in amplifiers.

2.4.2.2 Acoustic Signal Injection

Other research focused on using acoustic signals as a mechanism for signal injection. Several works [63, 64, 65, 66] showed that inertial sensors like MEMS accelerometers and gyroscopes can be jammed or spoofed by high-intensity acoustic signals near

resonant frequencies. Beyond just the inertial transducers, these works targeted vulnerable signal processing components such as amplifiers, filters, and ADCs to control the inertial sensor output. Other work showed that the systems relying on these inertial sensors could be affected even if sensor fusion techniques were in place [67]. *DolphinAttack* [68] and several other works [69, 70, 71] showed that microphones were vulnerable to signal injection via the demodulation of inaudible ultrasound. Some works demonstrated an acoustic spoofing attack on ultrasonic sensors used in automotive contexts [56, 72]. *Blue Note* [73] showed how drop sensors in certain hard drives were sensitive to acoustic signals, causing drops in throughput. Finally, *Poltergeist* [74] showed that acoustic signals could be used to influence image stabilization in cameras to cause captured images to be blurred.

2.4.2.3 Light-Based Signal Injection

Most related to my research, there have been several works investigating methods of injecting signals into sensors via light and lasers. Petit .et. al. [75] and Yan et. al. [76] investigated the use of lasers to perform a denial-of-service attack on cameras used in autonomous vehicles. From here numerous works of using light as a way to influence cameras and computer vision. *Phantom of the ADAS* [77] used a powerful projector to generate fake pedestrians and signs. *GhostImage* [78] used a projector to exploit lens flare effects within cameras to fool computer vision. Köhler et. al.[79] and *Rolling Collors* [80] used pulsed lasers to exploit the camera rolling shutter effect to fool computer vision. Sato et al. [81] showed how infrared lasers may be used to fool computer vision algorithms when cameras are lacking IR filters.

Beyond cameras, several other sensors have also been targeted by LSI. Park et al. [82] showed that medical infusion pumps could give incorrect doses if the IR drip sensor was injected with infrared light signals. Works such as Petit et. al. [75] and

Illusion and Dazzle [83] showed how infrared lasers can be used to inject false points into a LiDAR point cloud. Shin et. al. [84] showed how lasers could be used to spoof signals to light-based smoke detectors. Finally, Tanaka and Sugawara [85] used lasers to inject signals into piezoresistive MEMS pressure sensors.

2.4.2.4 Signal Injection Defenses

Though many works have focused on vulnerabilities and countermeasures for specific signal injection attacks, a few have attempted to develop generalized defenses to signal injection attacks. *PyCRA* [86] described a mechanism to detect injection within active sensors by selecting random times to stop active probing of the environment. If a strong signal is still captured without active probing, an adversarial attack can be inferred. This mechanism was applied to detect IEMI in passive sensors by Zhang and Rasmussen [87] by periodically powering down the sensor and monitoring the signal. Sun et. al. [88] showed how improving the sensor model from the software can lead to better anomaly detection in LiDAR spoofing attacks. Several other works have used various techniques and signal processing to detect anomalies within a stream of sensor data used in control systems and robotics [89, 90, 91].

CHAPTER III

Characterizing Laser Signal Injection on LiDAR

As the number of Autonomous Vehicles (AVs) on the roads increases [92, 93, 94], there is an increasing reliance on sensors to provide accurate information about the world around them. Among the sensors being employed by these vehicles is LiDAR, a complex, time-of-flight sensor used to develop three-dimensional mappings of the environment. While LiDAR provides information-rich data about the state of the system, it is vulnerable to laser signal injection.

The purpose of this chapter is to give an overview of laser signal injection attacks on LiDAR sensors, especially the ones used in autonomous vehicles. This work directly contributed to one publication[1], and enabled several follow-up works by other researchers. This chapter will answer the following questions:

- How are LiDAR sensors vulnerable to LSI?
- What capabilities does an attacker have in modifying LiDAR data with LSI?
- What are the consequences on systems that use LiDAR?
- How can future systems defend against LSI attacks on LiDAR?

[1] Y. Cao, C. Xiao, B. Cyr, Y. Zhou, W. Park, S. Rampazzi, Q. A. Chen, K. Fu, and Z. M. Mao, "Adversarial sensor attack on lidar-based perception in autonomous driving," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019, pp. 2267–2281.

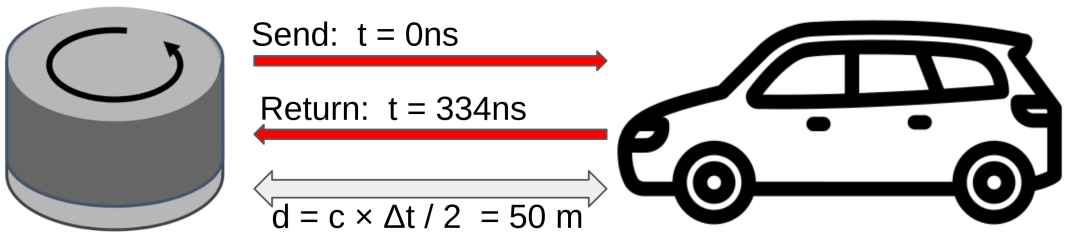

Figure 3.1: LiDAR functions by measuring the time-of-flight of an infrared laser beam. The time can be used to determine the distance to objects around the sensor.

3.1 LiDAR Sensors and their Applications

LiDAR functions by pulsing s low-power, infrared (IR) laser at the surrounding environment and capturing the reflection with s light-sensitive Avalanche Photodiode (APD). A reflection indicates the presence of an object. The amount of time between the laser pulse and when a reflection is captured can be used to calculate a distance (d) to the object (Fig. 3.1. This is calculated by:

$$d = c \times \Delta t/2 \tag{3.1}$$

where Δt is the time it takes between the laser firing and measuring a return, and c is the speed of light. A real-time 3D map of the surroundings is constructed when the IR lasers are pulsed and measured from many different azimuthal and altitudinal angles. This map is made up of many different points in space around the LiDAR, and the resulting points together make a point cloud.

With signal processing and machine learning algorithms, autonomous systems can be designed to use these point clouds to localize their position within the environment and avoid obstacles. This allows autonomous systems such as advanced robotics and autonomous vehicles to traverse through an uncertain environment with more information about the surroundings. Ultimately, this leads to devices that have less

uncertainty about their position and the position of potentially dangerous obstacles around them.

This work focuses on autonomous vehicles, as they are currently the most common use of advanced LiDAR systems. Autonomous vehicles have an inherently high potential for harm due to their mass, speed, and interactions with complex road environments. Because of this, LiDAR is seen as a useful tool to improve the safety and reliability of autonomous vehicles, reducing the likelihood of crashes. But the optical nature of LiDAR makes autonomous vehicles vulnerable to laser signal injection.

3.2 Related Work

Previous research work has shown that LiDAR systems are vulnerable to LSI using IR lasers. The APDs used in LiDAR rely on the photovoltaic effect (Sec. 2.3.1.2), and are extremely sensitive to even the smallest incoming irradiances. While this makes the LiDAR sensor much better at detecting the environment, it also means an adversary can inject a signal to the sensor when it is expecting a reflection from the environment.

This was first demonstrated as a replay attack by Petit et al. [75], where a laser was used to replay the IBEO LUX 2 LiDAR's light signal to inject a fake wall at distances greater than 40 meters from the LiDAR. This work was improved in *Illusion and Dazzle* [95], which presented two different attacks on a Velodyne VLP-16 LiDAR: saturation attacks and spoofing attacks. In the saturation attacks, a bright IR light is shined at the LiDAR, saturating the APDs and causing the system to miss true reflections. While the saturation attack is strong, it can be easily detected by anomaly detection techniques, as the resulting signal will be noisy or empty of reflections.

The stronger spoofing attack presented in *Illusion and Dazzle* was a spoofing attack. In their work, the precise firing of an IR laser could inject fake points in the LiDAR point cloud. This attack could be used to make it seem as if there were

objects that were reflecting light, even though no objects existed. Spoofing attacks are much harder to detect, as it doesn't hide the true reflections in the environment. Beyond this, the authors also showed that with precise timing, points could be spoofed between the spoofer and the LiDAR sensor. This presented a new threat model where an adversary could set up a laser at a long distance, yet spoof points very close to the LiDAR. This threat model was considerably stronger.

There were several limitations to this spoofing attack. One, the laser that was used had a limited firing rate. Because of this, the authors were only able to spoof around 10 points in a 2-degree azimuthal range, and a single altitudinal angle. Ten points, while still significant, are usually not enough for an object detection algorithm to recognize a point cloud as an obstacle. Beyond this, there was no way to shape the points in a way to exploit a machine-learning algorithm. The spoofed cloud would appear in the LiDAR's signal, but without any shape, it may only be recognized as noise. These limitations inhibited the attacker's capability to inject point clouds that could affect an autonomous vehicle system.

3.3 Laser Signal Injection Attacks on LiDAR

My work in this chapter further explored and characterized attacker capabilities in LSI attacks on LiDAR, with the purpose of showing the true level of vulnerability and begin highlighting defenses. In all of these investigations, the target LiDAR is a Velodyne VLP-16, which is a low-end, mechanically rotating LiDAR with 16 lasers. This sensor was used due to its low cost, but the principles these LSI attacks will apply to any LiDAR sensor that fires a sequence of lasers in a regular pattern.

3.3.1 Attack Overview

The principle of an LSI attack on LiDAR is to inject a laser signal that is perceived as a reflection of the environment. A spoofing attack consists of three primary

components:

1. **Photodiode**: A photodiode is needed to measure an infrared pulse from the lidar. The photodiode is necessary to synchronize the attack with the pulses of the victim lidar. Without the photodiode, the timing of pulses will be inaccurate, as there is no feedback mechanism to determine the current state of the LiDAR. In this setup, an OSRAM SFH 213 FA photodiode was used, with a simple circuit to bias the diode and increase its sensitivity.

2. **Delay Component**: Some type of delay or timing component is necessary to generate a series of precise pulses. This must be precise circuitry that can trigger laser pulses with nanosecond resolution. This can be a high-resolution function generator, FPGA, ASIC, or many other devices. This component is triggered by the photodiode, waits for a certain amount of time, and then sends one or multiple pulses to trigger the laser. In this project, an AFG3251 arbitrary function generator was used as the delay component. The AFG3251 has a timing resolution down to half a nanosecond, which gave increased capabilities to generate arbitrary firing patterns with very precise timing.

3. **Infrared Laser**: An infrared laser is necessary to spoof a reflection to the lidar. By sending out laser pulses with the correct timing, an attacker can create the illusion that objects have reflected light when no such objects exist. The laser used in this project was the OSRAM SPL PL90 laser diode, being driven by a PCO-7114 laser driver module. The laser driver is necessary to produce the correct current spikes to send to the laser diode, and the PCO-7114 has a maximum pulse repetition rate of 1.1 MHz, corresponding to a laser pulse every 900 ns.

A spoofing attack is illustrated in Figure 3.2. The attack works by using the photodiode circuit to measure when the LiDAR system has fired an IR laser in the

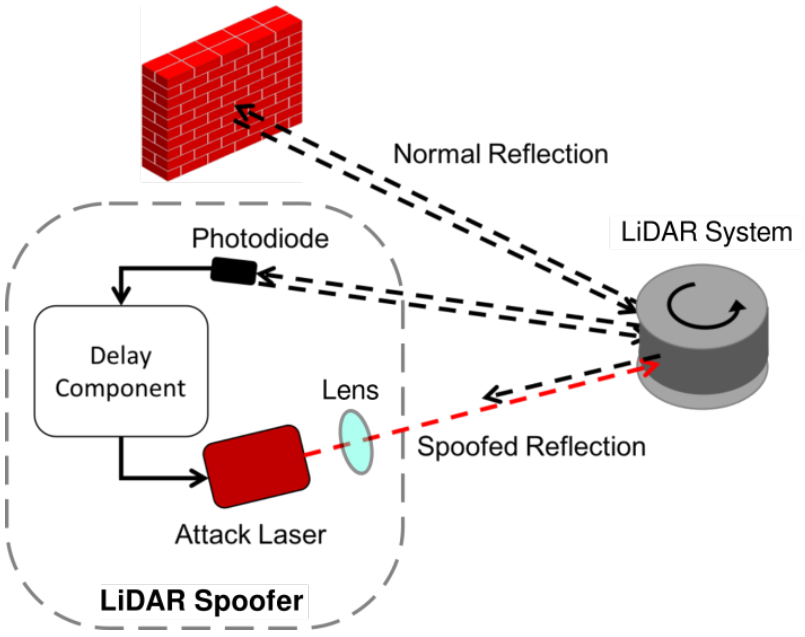

Figure 3.2: Illustration of LiDAR spoofing attack. The photodiode receives the laser pulses from the LiDAR and activates the delay component that triggers the attacker laser to simulate real echo pulses.

direction of the spoofer. Once the spoofer receives a signal from the LiDAR, it synchronizes with the timing of the laser firings. The spoofer then uses the delay component such as a function generator to wait a very precise amount of time before firing a signal back at the attacking laser. By changing the delay, the attacker can change the distance of the spoofed point from the LiDAR, even injecting a fake point cloud between the LiDAR and the spoofer. This has much greater consequences, as it allows the attacker to set a laser far away and still generate fake obstacles in regions that are close to the vehicle.

3.3.2 Characterizing Attacker Capabilities

While previous work had shown a spoofing attack was shown to be possible, there was very little effort in exploring and characterizing an attacker's full capabilities. Several improvements to the attack setup in previous works enabled a better definition of the attacker's capabilities. These improvements and capabilities are summarized

here.

3.3.2.1 Number of Points

Improvements to the laser driver triggering enabled the attacking laser to fire at a faster rate. Previous works only fired in steps of 110 μs [95], but a faster-firing laser on the order of microseconds gives an attacker significantly more control over the victim LiDAR. With the increased speed and control offered by the laser driver and arbitrary function generator, spoofed points could be generated at all 16 altitudinal angles on the VLP-16 LiDAR, with up to 100 spoofed points or more. An example of a large point cloud is shown in Figure 3.3, where the cloud extends to multiple vertical angles

The only limitations to the number of points are the speed at which the attacker can fire lasers and the horizontal field of view of the LiDAR. The VLP-16 is relatively slow in the world of advanced LiDAR systems, firing once every 2.304 microseconds. Our laser setup had a maximum firing rate of 900 nanoseconds, well within the firing period of the VLP-16. This may not be the case with other LiDAR systems. Beyond this, the setup could only spoof points within an 8° horizontal angle. Outside of this angle, the injected laser light would not be focused properly onto the APDs by the LiDAR optics, so no signal would be measured. A potential way to overcome this limitation is to compensate with a higher-power injection laser. By exploiting reflections and lens flare effects within the optics of the LiDAR sensor, it would be possible to extend this horizontal range and increase the number of points. Unfortunately, due to safety concerns, this setup was not capable of reaching these experiments and must be left to future work.

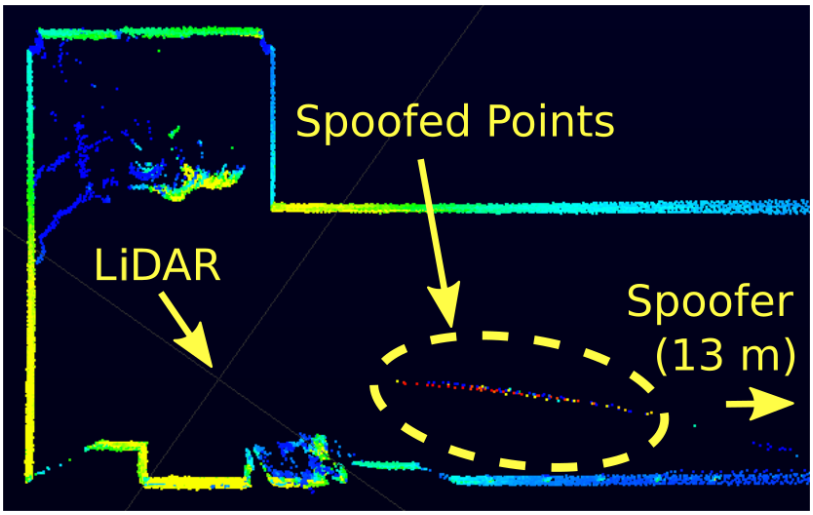

Figure 3.3: Collected traces from the reproduced sensor attack. The large number of points in the yellow circle is spoofed by the sensor attack at multiple altitudinal angles.

3.3.2.2 Range

Another improvement was the addition of focusing optics to greatly extend the attack range by increasing the irradiance. Infrared lasers diverge quickly without a collimating lens, which was limiting the effectiveness of the attack in previous works. To show the increased range provided by the optics that were added to the attacking laser, a long-distance spoofing experiment was performed on the VLP-16. The experiment was performed with a simple lens in a 15-meter-long windowless hallway. With this setup, large point clouds between the spoofer and the target were spoofed with the spoofer up to 13 meters away, as shown in Figure 3.3.

This 13-meter range is on the very low end of what is possible, as the 15-meter hallway was a limitation due to safety concerns. From the model of beam propagation discussed in Section 2.2.3, we know that significant irradiances can be sent over long distances with special optical equipment to expand the beam of the laser. Something as simple as a telephoto lens could greatly increase the range. The APDs in LiDAR sensors are designed to be extremely sensitive to even small irradiances. With better optical equipment and calibration, it seems feasible that this range could extend to

well over a hundred meters. More work would need to be done to determine what irradiances would be required.

3.3.2.3 Point Shaping

The most important improvement to the spoofing attack was the addition of an arbitrary function generator to enable the shaping of an injected point cloud. A function generator could be used to generate an arbitrary firing sequence, rather than only pulsing at a single frequency. With the features to generate arbitrary waveforms, a higher level of control was available in generating the attacking laser pulses. With this addition, there was a shift in the strategy of the attack. Previous works focused on generating as many spoofed points as possible. With this increased capability, the strategy shifted away from generating as many points as possible to shaping the spoofed point cloud in a way to best exploit an object detection algorithm. This is done by selectively choosing when to fire laser pulses at the LiDAR using a pattern programmed into the arbitrary waveform generator.

To illustrate the change in attack strategy towards shaping point clouds, a spoofed point cloud is shown in Figure 3.4. In this case, the attacking signal can be used to spoof points in the shape of "UM", showing that an attacker has the capability to control the spoofed point cloud and perform various transformations upon it.

The only limitation of the point cloud shaping was the capabilities of the arbitrary waveform generator. It is difficult to have nanosecond-level precision to precisely spoof a point cloud. The memory limitations of the function generator in this work meant only a small section of the total point cloud could be shaped. This limitation could potentially be overcome with more modern equipment or specially-engineered devices with FPGAs or ASICs that can generate precise and programmable firing signals.

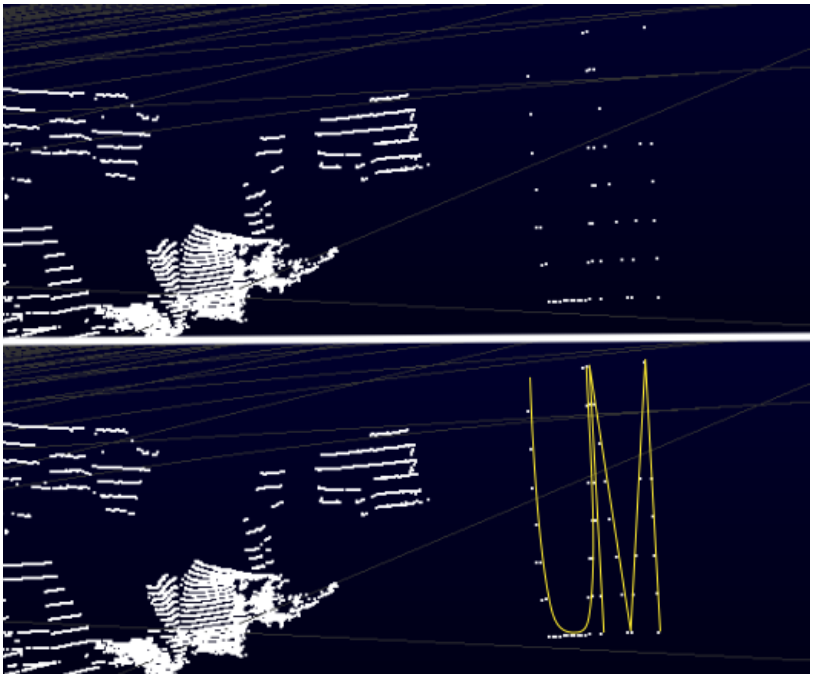

Figure 3.4: Since the VLP-16 firing sequence is predictable, my experiments showed the spoofed point cloud can be shaped by the attacker.

3.3.3 A Model of LSI in LiDAR

The ultimate goal of an LSI attack on LiDAR is the injection of a set of points that will be interpreted as an obstacle. To do this, an attack not only has to inject a significant number of points, but shape them in a way to cause the underlying algorithms to perceive the injected point cloud as an obstacle. To generate this shape, an attacker needs to generate a firing sequence with the attacking laser. This firing sequence must have very precise timing, as the speed of light ensures even slight delays cause large variations in the spoofed point distance.

The process of generating this firing sequence can be modeled in the following way. Suppose the attacker has an amount m of points in a point cloud P that will be injected into the LiDAR's measured point cloud.

$$P = \{(d_1, \theta_1, \phi_1), (d_2, \theta_2, \phi_2), \ldots (d_m, \theta_m, \phi_m)\} \tag{3.2}$$

Since the point cloud is three-dimensional, each injected point will have three coordinates relative to the body frame of the LiDAR. This is shown by the spherical coordinates d (the distance from the LiDAR to the point), ϕ (the azimuthal or horizontal angle), and θ (the altitudinal or vertical angle).

To inject this point cloud, these points need to be converted into a laser-firing sequence. This sequence can be shown by a set of delays T_0, that describe the delay between measuring the LiDAR signal and firing a laser back.

$$T_0 = \{\tau_1, \tau_2, \tau_3, \ldots \tau_m\} \tag{3.3}$$

Each delay τ_i corresponds to a single point to be injected into the point cloud, and can be described by:

$$\tau_i = \tau(d_i, \theta_i, \phi_i) - \tau(d_0, \theta_0, \phi_0) - \tau_d \tag{3.4}$$

where d_0, θ_0, and ϕ_0 are the coordinates of the spoofer, and τ_d is the inherent delay of the attacking hardware. The delay timing function τ is described by:

$$\tau(d, \theta, \phi) = \frac{2d}{c} + f_\theta(\theta)\tau_\theta + f_\phi(\phi)\tau_\phi \tag{3.5}$$

where c is the speed of light, τ_θ is the LiDAR characteristic altitudinal delay, and τ_ϕ is the LiDAR characteristic azimuthal delay. These characteristic delays are properties of the LiDAR and are determined from documents or experiments on the target LiDAR device. The mapping functions f_ϕ and f_θ map the desired angles to non-negative integer values:

$$f_\phi(\phi) \in \mathbb{N}_0 \qquad , \qquad f_\theta(\theta) \in \mathbb{N}_0 \tag{3.6}$$

where $\mathbb{N}_0$ is the set of natural numbers including zero. These mapping functions

are also unique to each LiDAR and need to be determined from documentation or experimentation on the target device.

The set of delays T_0 will describe a firing sequence to inject the point cloud data a single time, but since the LiDAR constantly measures new signals, the pattern needs to be constantly repeated, as shown by:

$$T_i = \{\tau_1 + i\tau_R,\ \tau_2 + i\tau_R,\ \tau_3 + i\tau_R, \dots\ \tau_m + i\tau_R\} \tag{3.7}$$

where τ_R is the period of time the LiDAR uses to collect a single scan. In theory, this allows an attacker to inject a point cloud indefinitely after collecting a single laser pulse from the LiDAR. In reality, timing errors will accumulate, requiring the attacking laser to be regularly resynchronized with a new pulse from the target device.

From this model, we can see that the true secrets being exploited by the LSI is the mapping functions and the characteristic delays of the LiDAR. As an example with the Velodyne VLP-16, knowing the properties of the LiDAR allows the attacker to inject and adjust a large range of spoofed point clouds. Research into the VLP-16 device yielded that the characteristic delays were $\tau_\phi = 55.296\ \mu s$ and $\tau_\theta = 2.304\ \mu s$. This is found both experimentally and by a view of the documentation. The overall scan period τ_R is configurable, but by default, it is $\tau_R = 100\ ms$ as the rotational frequency is set to $10\ Hz$. The mapping functions are shown by:

$$f_\phi(\phi) = \left\lfloor \frac{(\phi - \phi_0 + 0.1^\circ) \mod 360^\circ}{0.2^\circ} \right\rfloor \tag{3.8}$$

$$f_\theta(\theta) : \{-15^\circ, 1^\circ, -13^\circ, 3^\circ, \dots, -1^\circ, 15^\circ\} \mapsto \{0, 1, 2, 3, \dots, 14, 15\} \tag{3.9}$$

demonstrating the limitations on the angles that can be chosen in order to map to the set of non-negative integers. Using this information, an example of the firing signal

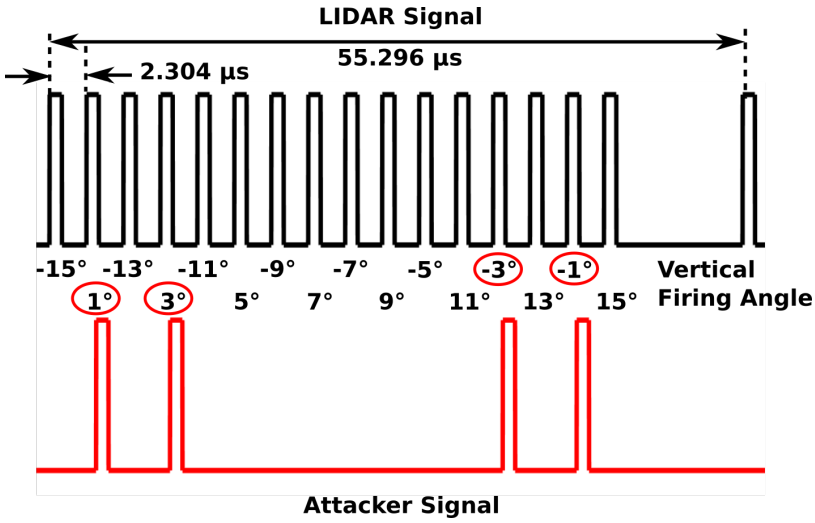

Figure 3.5: The consistent firing sequence of the LiDAR allows an attacker to choose the angles and distances from which spoofed points appear. For example, applying the attacker signal, fake dots will appear at 1°, 3°, -3°, and -1° altitudinal angles.

is shown in Figure 3.5, showing how an attacker can selectively spoof points at only four altitudinal angles.

This model leads to a better understanding of what an attacker can do while performing a LiDAR spoofing attack. Beyond just making false point clouds appear, an attacker has the capability to transform the spoofed point cloud in three different ways. These capabilities are shown in Figure 3.6. First, if the spoofing delay is changed in small steps on the order of nanoseconds, the d coordinate of the spoofed point can be changed (Fig 3.6a). For every nanosecond change, the point moves by around 15 centimeters. Second, if the spoofing delay is changed in steps of 2.304 microseconds, there is a change in the altitudinal angle θ of the spoofed point, allowing an attacker to change the vertical position of the point (Fig 3.6b). Finally, the azimuthal angle ϕ can be changed by changing the delay in steps of 55.296 microseconds, changing the horizontal position of the point (Fig 3.6c). With these capabilities well-defined, an end-to-end attack on autonomous vehicle LiDAR could be developed.

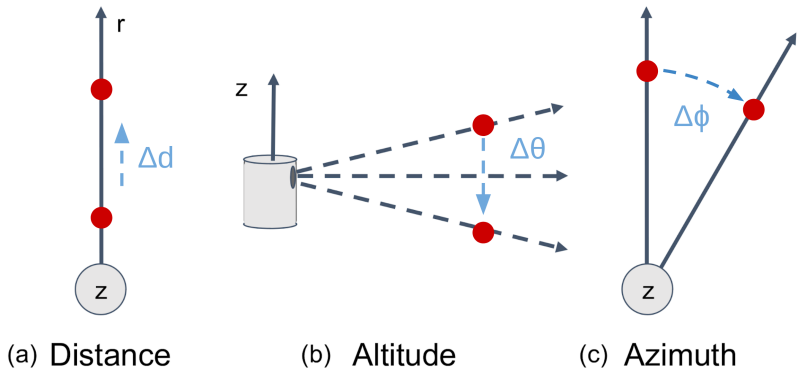

Figure 3.6: Attacker capabilities in spoofing points in a point cloud by changing the timing of laser firings

3.4 Consequences on Autonomous Vehicles

The exploration and definition of the attacker capabilities enabled a collaborative work that was published at the 2019 ACM CCS [4]. That paper was focused on generating adversarial spoofed point clouds that would be classified as objects in object detection algorithms. The algorithm for generating these point clouds is dependent on the model of attacking capabilities that were defined by these experiments. While the attack scenarios in this work were limited to vehicles that are stationary relative to the attacking laser, it showed that there is a vulnerability in LiDAR that can potentially have life-threatening consequences.

In particular, the paper used a simulation of the Baidu Apollo vehicle system to propose two attacks: a freezing attack and an emergency braking attack. In the freezing attack, laser signal injection is performed on a stopped vehicle, causing a false object to appear in front of the vehicle. The control system does not know how to handle the false object and therefore freezes while it waits for the object to disappear. This can cause traffic build-up and potentially be threatening to incoming traffic. The braking attack is performed by mounting an attack laser on one vehicle and injecting a false object in front of another vehicle while traveling at high speed. The vehicle assumes that the false object is an obstacle and performs an emergency

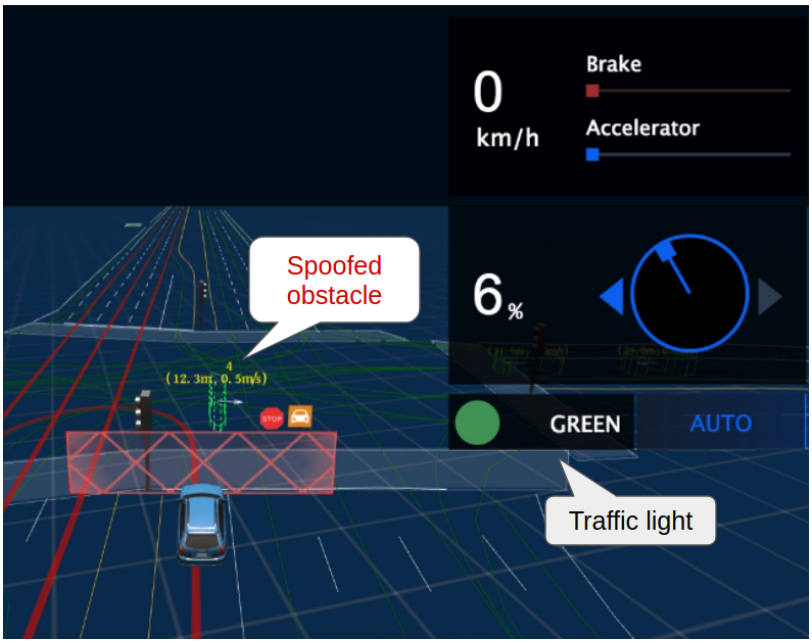

Figure 3.7: A freezing attack performed against an AV. An obstacle spoofed in an intersection prevents the AV from progressing through the intersection

stop, threatening the lives of the passengers and anyone else on the road behind them. These are just two simulated consequences on a single autonomous vehicle system, but there are many more possibilities.

It is difficult to fully explore the true consequences of LSI on autonomous vehicles in general. Every AV will have differences in signal processing and algorithms that make it difficult to predict the effectiveness of an LSI attack. Even slight adjustments to the algorithms may make it much more difficult to inject a spoofed obstacle. Since these algorithms will be a black box to most attackers, the effectiveness of any real-world LSI on LiDAR will depend on the information known and the assumptions made by the attacker. Even if the principles of laser signal injection are established, the effects on other AVs must be left to future work.

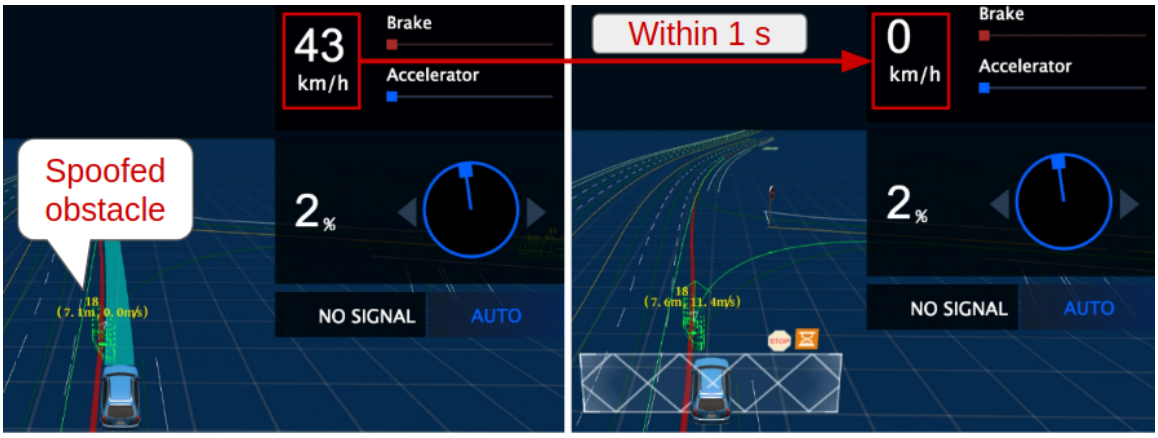

Figure 3.8: An emergency brake attack performed against an AV. An obstacle spoofed in front of a vehicle on a highway causes it to initiate an emergency brake

3.5 Future Directions

There is still much to be done to research the effects of LSI on LiDAR sensors and autonomous vehicles. This includes an investigation of various potential defenses, as well as several open problems that still exist within the AV space.

3.5.1 Recommendations for Defenses

The experiments and models developed for laser signal injection in LiDAR have shown the capabilities for an attacker to have significant control over LiDAR data, with many potential consequences on the target system. Future devices should be designed in a way to reduce or remove these capabilities. Many defenses have been suggested to prevent these attacks. These can be detection mechanisms or mitigation techniques and generally fall into two categories: system-level defenses and sensor-level defenses.

3.5.1.1 System-Level Defenses

System-level defenses require changes to the autonomous vehicle system in the form of extra components or updates to algorithms. These defenses are implemented by the designer of the AV.

An unintentional defense already present in most AV systems is the use of sensor fusion techniques to combine data from multiple or redundant sensors to determine if sensor data is faulty. If data from a single sensor detects an obstacle, yet two other sensors monitoring the same part of the environment do not detect an obstacle, then it may be an indication of an attack. Sensor fusion is intentionally included for the purpose of reliability, as data from multiple sensors will give more information about the environment. In the case of an attack, however, it is difficult to develop a robust strategy in the event of an anomaly between sensors. Because autonomous vehicles are high-risk systems, sensor fusion techniques tend to be cautious, leaving the system more vulnerable to attacks that cause slowdowns and braking. Beyond this, sensor fusion inherently requires the collection of redundant information about the environment. As technology develops, the expected trend would be to reduce costs by lowering the number of sensors collecting redundant information, reducing the effectiveness of sensor fusion. Future work is needed to develop strategies that can handle adversarial signals, not just environmental noise.

Another area to investigate potential defenses is in the object detection algorithms that rely on LiDAR data. In our work [4], it was discovered that the Baidu Apollo object detection algorithm did not incorporate altitudinal information from the LiDAR data, making it much easier to perform a spoofing attack by using ground reflections in the environment. Improvements to these algorithms and the incorporation of extra temporal and spatial context into the object detection algorithms would help detect anomalies within the point cloud that could indicate an attack. Research into this area has led to a follow-up paper discussing some potential updates to these

algorithms [88], but more work should be done to evaluate the effectiveness of these changes.

3.5.1.2 Sensor-Level Defenses

Sensor-level detection defenses are changes to the sensor hardware to detect the presence of LSI. Often these defenses will require significant engineering effort and advanced knowledge of the LiDAR hardware, and can only be implemented by the sensor designers.

LiDAR sensors can be improved by detecting an incoming laser signal injection. Perhaps the simplest mechanism is the solution proposed by PyCRA [86], where the LiDAR skips a laser firing with some known but random frequency. The idea of PyCRA is that if the sensor senses a return when no laser was fired, it would indicate the presence of an attack. This inherently comes with a tradeoff. For every laser firing that was skipped, it lowers the resolution of the data, which could potentially harm the performance of the AV system. Another potential detection mechanism would be to look for anomalies in the APD signal. A measured laser pulse signal will have certain characteristics, which could potentially be used to identify a laser pulse from the sensor versus a spoofing laser.

Other Sensor-level defenses can potentially mitigate LSI against LiDAR sensors. As mentioned in Section 3.3.3, the primary secret being exploited by an attacker is the timing of the laser pulses and the mapping of that timing to each angular coordinate in the collected point cloud. Randomized sampling techniques can be used to better hide these secrets by adding randomness to the firing rate or pattern. While this would definitely make LSI significantly more complex, more research is required to determine the appropriate amount of randomness to defend against a spoofing attack yet still provide a similar amount of information about the environment as a consistently scanning LiDAR

Another defense is to change the optical components present within LiDAR systems. LiDAR sensors were originally designed with wide-view optics to capture as much light from the return as possible. This lead to large lenses and the minimizing of optical barriers between the reflected signal and the sensitive APDs. While improving the performance of the LiDAR, these optical conditions allow an attacker to affect a wide field-of-view by exploiting lens flare and other optical effects that bend incoming light to the incorrect APD. A single device can then be used to spoof points over a wide range of the measure point cloud. With some changes to the optics to restrict the light to travel only the expected optical paths, the vulnerable field-of-view could be reduced.

Finally, a mitigation technique has already been developed by encoding a signal onto the laser pulses as they are fired from the LiDAR. Some companies have developed systems that will encode a pattern onto each laser signal as they are fired to measure the environment [96]. Each pattern corresponds to a single laser firing and can be used to verify the return. This encoding system was developed to improve reliability and reduce false reflections, but it also greatly increases the complexity of a spoofing attack. More research is needed to determine the effectiveness of this defense.

3.5.2 Limitations and Open Problems

While this work sought to characterize the current capabilities of LSI in LiDAR sensors, there is many open problems regarding the limits and consequences of this attack.

A primary limitation of this attack is the problem of aiming. While systems can be designed with very precise timing, it is more difficult to design a system to consistently track a target and understand its relative position to spoof a consistent point cloud. Significant engineering effort will be needed to design a system that

can handle the motion of the target. Without this, the attack is limited to targets stationary relative to the spoofer.

Another limitation of this LSI attack is that it can only add points to a spoofed point cloud. Any true reflections will still be captured. As jamming attacks in previous works have shown LiDAR to be vulnerable to denial-of-service, this opens up the possibility of selective hiding attacks. By firing a laser with precise timing, the true signal can be removed, which will also have considerable consequences on the security of autonomous vehicles. This has already been explored by one follow-up work [97], but there are open questions regarding how these two classes of attacks can be combined to better understand the attacker's capabilities.

Other limitations of the current setup are the range and speed of laser firing. The optics used in this setup were simple, with only a small space to perform our experiments. There are many ways in which future works can greatly increase the range of these attacks. Beyond this, with increases in irradiance, there may be a greater capability to influence a large field-of-view by exploiting lens flares and reflections within the LiDAR optics itself. Beyond range, many modern lasers are being developed with more lasers that fire more quickly. As the speed of the lasers increases, there will need to be further experiments to determine the capabilities of attackers with various firing rates of spoofing lasers. Future work is required to fully determine the extent of these limitations.

Significant work remains in determining the real-world consequences of laser signal injection on autonomous vehicles. While this research is expanding and characterizing attacker capabilities, the costs of determining the effects on real systems are considerable. There are many companies developing new AV systems with black box techniques. Future work is needed to provide test procedures, simulation tools, and data sets to determine the true consequences of laser signal injection in vulnerable LiDAR sensors.

CHAPTER IV

Characterizing Laser Signal Injection on MEMS Microphones

The consistent growth in computational power is profoundly changing the way that humans and computers interact. Moving away from traditional interfaces like keyboards and mice, in recent years computers have become sufficiently powerful to understand and process human speech. Recognizing the potential of quick and natural human-computer interaction, several companies have launched their own large-scale deployment of low-cost Voice-Controllable Systems (VCSs) that continuously listen to and act on human voice commands.

The ubiquity of VCSs is made possible by the development of the MEMS microphone, a small and low-cost way to sense acoustic pressure signals in the air. This sensor has become a common addition to many IoT devices to enable new features and capabilities. Recently, however, it was discovered that MEMS microphones are vulnerable to LSI. The work in this chapter [1] [2] [3] demonstrates how LSI can be used

[1] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, "Light Commands: Laser-Based Audio Injection Attacks on Voice-Controllable Systems," in *29th USENIX Security Symposium (USENIX Security 20)*, 2020, pp. 2631–2648.

[2] B. Cyr, T. Sugawara, and K. Fu, "Why Lasers Inject Perceived Sound Into MEMS Microphones: Indications and Contraindications of Photoacoustic and Photoelectric Effects," in *2021 IEEE Sensors*, Oct. 2021

[3] B. Cyr, V. Sumaria, S. Tadigadapa, T. Sugawara, and K. Fu. "How Lasers Exploit Photoacoustic and Photoelectric Phenomena to Inject Signals into MEMS Microphones". Unpublished manuscript. April 2023.

to inject false acoustic signals into many vulnerable systems. This vulnerability was unexpected, and LSI was well outside the capabilities considered in previous threat models against voice-controllable devices. Beyond this, it was unclear what physical effects were being exploited.

The purpose of this chapter is to give an overview and investigation of laser signal injection attacks on MEMS microphones, especially the ones used in voice-controllable systems. It will answer the following questions:

- How can Micro-electro-mechanical Systems (MEMS) microphones be exploited by LSI?
- What capabilities does an attacker have in generating a false acoustic signal?
- What are the consequences on systems that use MEMS microphones?
- How can future systems defend against LSI attacks on MEMS microphones?

4.1 MEMS Microphones and their Applications

Microelectromechanical Systems (MEMS) are devices that consist of tiny structures that change their electrical properties as they experience mechanical movement. By measuring these electrical properties with an Application-Specific Integrated Circuit (ASIC), MEMS devices output an electrical signal that corresponds to mechanical motion. In particular, MEMS microphones sense mechanical vibrations of a thin plate (called a diaphragm) that responds to acoustic energy, outputting an electrical signal that represents the sound hitting the microphone. Figure 4.1 shows the typical construction of one of these microphones, where an open port is present in the microphone to allow acoustic waves to hit the diaphragm. As the diaphragm vibrates, there is a change in capacitance between the diaphragm and a stationary backplate, which is measured by the ASIC. After this signal is filtered and amplified,

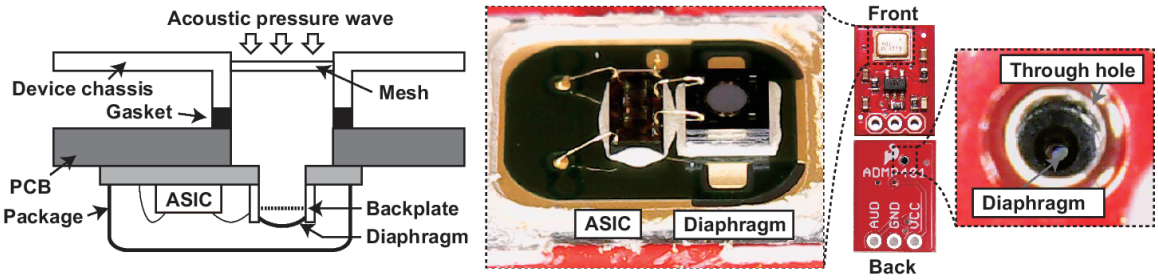

Figure 4.1: MEMS microphone construction. (Left) Cross-sectional view of a MEMS microphone on a device. (Middle) A diaphragm and ASIC on a depackaged microphone. (Right) Magnified view of an acoustic port on PCB.

the signal representing the sound is sent out to the rest of the system to be used in the appropriate application.

One important application of MEMS microphones is in Voice-Controllable Systems. Voice-Controllable Systems (VCSs) are devices such as smart speakers and smartphones that allow a user to interface with a computing system by speaking a natural language. These systems often include a programmed voice assistant to aid the user interface. Some examples of VCSs and assistants are iPhones with Siri, Google Home with Google Assistant, and Amazon Echo devices with Alexa. As these devices have become more common, there has been an effort to enable VCSs to complete as many tasks as possible, including tasks that should require special privileges. Some devices have enabled users to make online purchases, unlock smart locks, or open garage doors, all with the power of their voice.

4.2 Related Work

Because of the increasing popularity of VCSs, an entire line of research has developed around acoustic signal injection into these devices. The first attacks simply used malicious applications to play voice commands through a phone speaker to activate nearby VCSs [98, 99]. Because many of the voice commands do not require any authentication or privileges, the malicious commands are carried out. Further works

showed that commands could be hidden within audible signals so that an attacker could use VCSs in a way that a human could not recognize [100, 101, 102]. This was related to a line of work called skill squatting attacks [103, 104], which exploited errors in the recognition of similar-sounding words to install malicious applications on VCSs. Later works focused on creating inaudible commands to inject commands without alerting a nearby human. Roy et al.[105] showed that nonlinearities within microphones could be exploited to inject signals with ultrasound. Building off this, Song and Mittal [106] and *Dolphin Attack* [107] showed that commands could be injected into VCSs without any audible signal to alert a nearby user. This research was limited to a range of 1.75 meters but was extended by Roy et al. [108] and Yan et al. [71] to a maximum range of 19.8 meters using an array of ultrasonic speakers. While these attacks were practical and showed the vulnerabilities of VCSs, they were limited by sound attenuation due to air and obstacles such as windows.

While these VCSs are inherently vulnerable to signal injection with sound, there is some work showing that they are vulnerable to electromagnetic radiation as well. A few works [109, 110] showed that electromagnetic interference can be injected with radio frequencies on a microphone cord or even a power cord connected to a smartphone. By modulating a voice signal onto the electromagnetic signal, the authors were able to activate and inject voice commands into the VCS built into the smartphone. While building on past research to inject commands into these devices without being detected, the electromagnetic field that was required to inject a command was approaching the limit of human safety and well beyond the required immunity level of the device. This reduced the practicality of the attack, as it meant it would be difficult to perform the attack at range.

All of this research showed that VCSs were vulnerable to signal injection attacks, but all had limitations due to the physical medium with which the attack was being performed. This work focused on a different injection mechanism: laser signal

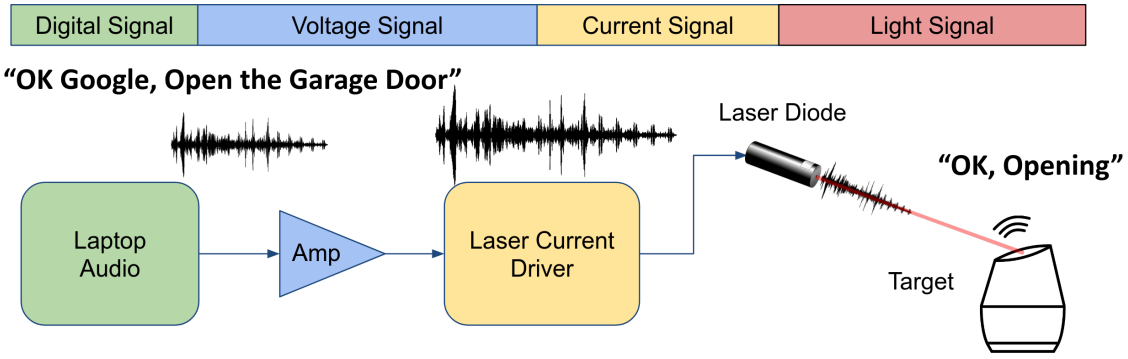

Figure 4.2: An overview of an LSI attack on MEMS microphones. An audio signal is converted by various components into an optical irradiance signal, which is measured by the microphone.

injection.

4.3 LSI Attacks on MEMS Microphones

Lasers can be used to inject a signal into MEMS microphones. This was a very surprising result to the security community, as no one considered the physical mechanisms that can be exploited to inject a signal with light. This led to several works with the purpose of investigating the causality, capabilities, and consequences of this new threat model.

4.3.1 Attack Overview

The fundamental principle in laser signal injection into MEMS microphones is that an amplitude-modulated laser irradiance signal will generate a voltage signal on the output of the microphone. An overview of the attack is shown in Figure 4.2. The attack requires 3 main components:

1. **Signal Source**: The desired acoustic signal must be generated as a voltage signal. This is most easily done with a laptop or phone that has the capability

of driving an audio voltage signal. This signal can be optionally filtered or amplified by external components to adjust the amplitude.

2. **Laser Driver**: A laser driver is a special component that converts the incoming voltage signal into a stable current signal to drive a laser diode. As mentioned in Section 2.2.2, the optical output power of a laser diode is proportional to the current across it. A laser driver ensures a linear transformation of voltage to current to generate the optical signal.

3. **Laser Diode**: A laser diode generates the irradiance signal required to interact with the target device. Many different colors of light can be used, as is explored later. Optics are required to shape the beam and extend the range, as described in Section 2.2.3. As this irradiance signal interacts with the target device, several possible mechanisms induce a voltage on the microphone output, which is interpreted as an acoustic signal.

As an example, a demonstration of LSI in MEMS microphones is shown in Figure 4.3. In the demonstration, a blue Osram PLT5 450B 450-nm laser diode is connected to a Thorlabs LDC205C laser driver. We increased the diode's DC current until it emitted a continuous 5 mW laser beam while measuring light intensity using the Thorlabs S121C photo-diode power sensor. The beam was subsequently directed to the acoustic port on a SparkFun MEMS microphone breakout board mounting an Analog Devices ADMP401 microphone. Finally, we recorded the diode current and the microphone's output using a Tektronix MSO5204 oscilloscope.

To convert sound signals into light, the voltage signal is encoded as an optical irradiance signal, where louder sounds make for larger changes in light intensity and weaker sounds correspond to smaller changes. Next, as the intensity of the light beam emitted from the laser diode is directly proportional to the supplied current, we use a laser driver to regulate the laser diode's current as a function of an audio

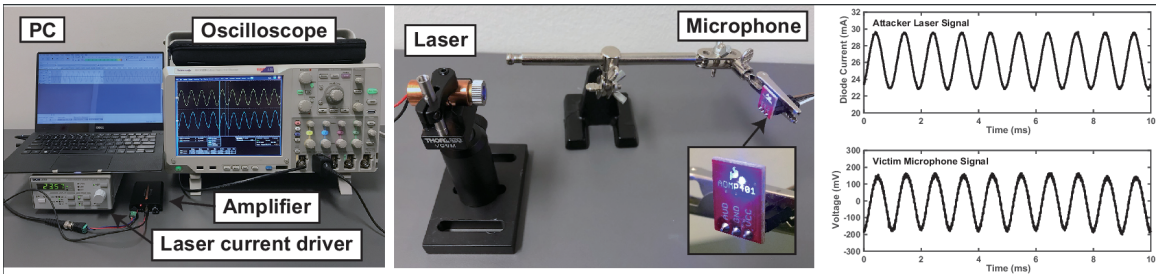

Figure 4.3: A demonstration of LSI on MEMS microphones. (Left) A setup for signal injection composed of a laser current driver, PC, audio amplifier, and oscilloscope. (Middle) Laser diode with beam aimed at a MEMS microphone breakout board. (Right) Diode current and microphone output waveforms.

file played into the driver's input port. This resulted in the audio waveform being directly encoded in the intensity of the light emitted by the laser.

More specifically, we used the current driver to modulate a sine wave on top of the diode's current I_t via amplitude modulation (AM), given by the following equation:

$$I_t = I_{DC} + \frac{I_{pp}}{2}\sin(2\pi f t) \tag{4.1}$$

where I_{DC} is a DC bias, I_{pp} is the peak-to-peak amplitude, and f is the frequency. In this section, we set $I_{DC} = 26.2$ mA, $I_{pp} = 7$ mA and $f = 1$ kHz. The sine wave was played using an onboard soundcard in a laptop, where the speaker output was connected to the modulation input port on the laser driver via a Neoteck NTK059 audio amplifier. As the light intensity emitted by the laser diode is directly proportional to the current provided by the laser driver, this resulted in a 1 kHz sine wave directly encoded in the intensity of the light emitted by the laser diode.

As can be seen in Figure 4.3, the microphone output clearly shows a 1 kHz sine wave that matches the frequency of the injected signal, with barely any noticeable distortion.

4.3.2 Characterizing Attacker Capabilities on VCSs

Once the vulnerability in MEMS microphones was discovered, it lead to the development of *Light Commands* [5], an LSI attack to inject commands into VCSs. Several experiments were performed to characterize the vulnerability within VCSs to determine the limits and capabilities of an attack.

Experiments were performed with the experimental setup shown in Figure 4.4, where a laser was fired at various VCSs to discover their vulnerability to light signal injection. Using a set of scanning mirrors, the laser beam was precisely aimed into the microphone ports to discover the minimum power needed to successfully inject commands into each of the devices. Seventeen different VCSs were chosen for the experiments, which are listed in Table 4.1. Four devices are smartphones and tablets, with limited voice authentication capabilities. For the experiments, we trained this authentication feature with the same voice that was later used in the signal injection. We defined a successful injection to be when a set of four distinct commands was correctly interpreted by the device three times in a row. These experiments lead to a better understanding of the differences between the VCSs when they respond to the light signal injection.

Next, experiments were performed to measure the effectiveness of the light injection at range. Figure 4.5 shows the setup of these experiments. The maximum effective range of the light signal injection was empirically measured on the seventeen VCSs at two different optical power levels: 5mW and 60mW. The 5mW power represents a low-power attack, with optical light equivalent to a laser pointer. The 60mW power represents the maximum power attack that could safely be performed with our equipment and setup. The experiments were performed in two different corridors: one with a length of 110 meters where the 5mW experiments could safely be performed, and one at a length of 50 meters where the 60mW experiments could safely be performed. A 650-1300mm telephoto lens (86mm diameter) was used to further

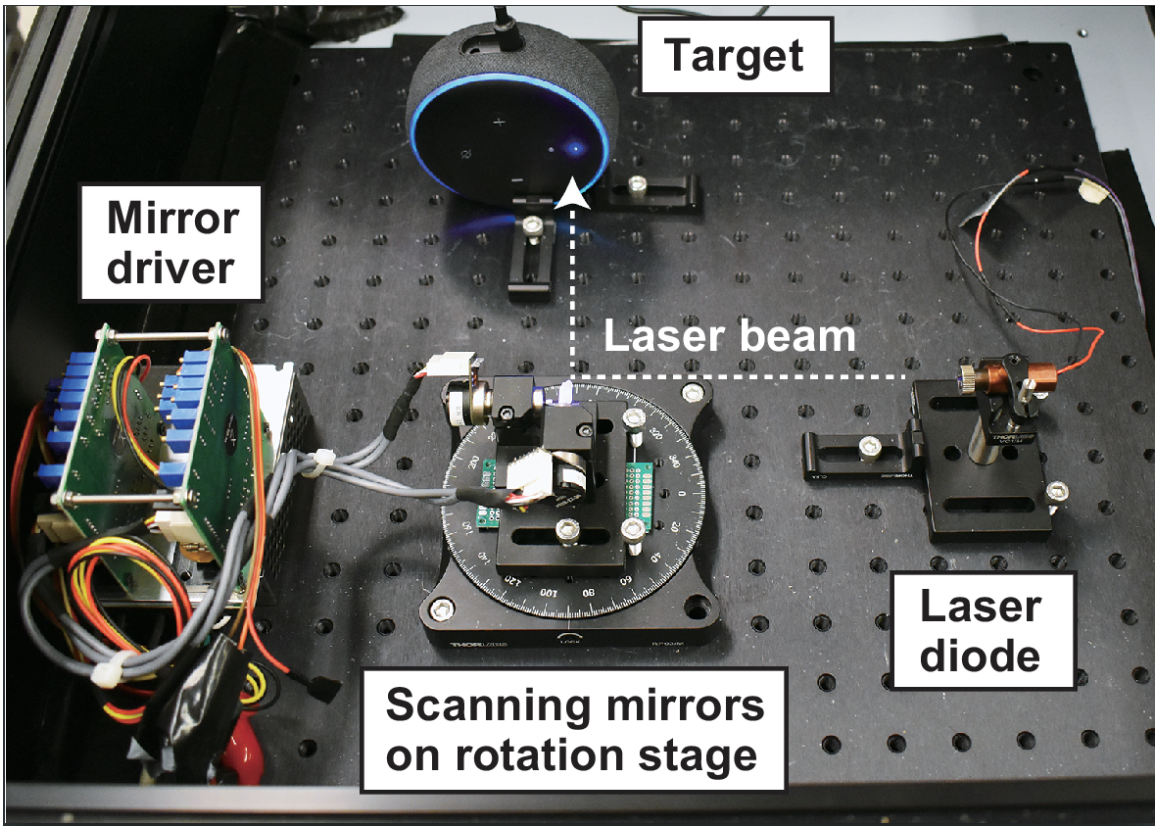

Figure 4.4: Setup for exploring minimum laser power requirements: the laser and target are arranged in the laser enclosure. The laser spot is aimed at the target acoustic port using electrically controllable scanning mirrors inside the enclosure. The enclosure's top red acrylic cover was removed for visual clarity.

focus the laser beam at longer distances, as this greatly increased the effectiveness of the attack at range. Just like in the enclosure experiments, a successful injection was defined to be when a set of four distinct commands was correctly interpreted by the device three times in a row.

Table 4.1 summarizes the results of the experiments on voice-controllable systems. The table shows that many smart home devices are especially vulnerable to Light Commands, with five devices having a minimum successful injection power of below 5mW, which is below the optical power of a standard laser pointer. Out of these five devices, the Google Home and Echo Plus 1st Generation were vulnerable even at our maximum range of 110 meters for the 5mW experiments. The majority of the other

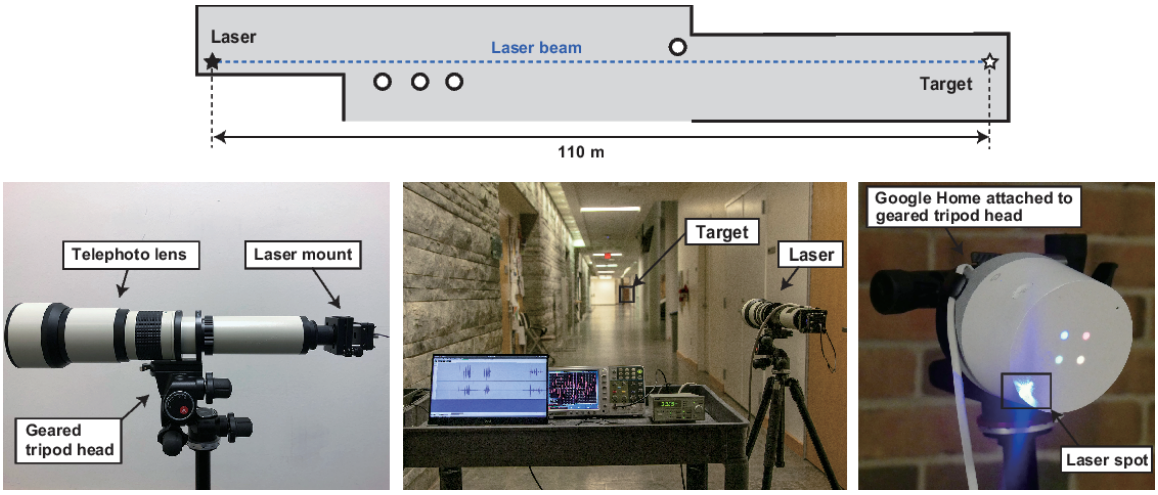

Figure 4.5: Experimental setup for exploring attack range. (Top) Floor plan of the 110 m long corridor. (Left) Laser with a telephoto lens mounted on geared tripod head for aiming. (Center) Laser aiming at the target across the 110 m corridor. (Right) Laser spot on the target device mounted on a tripod.

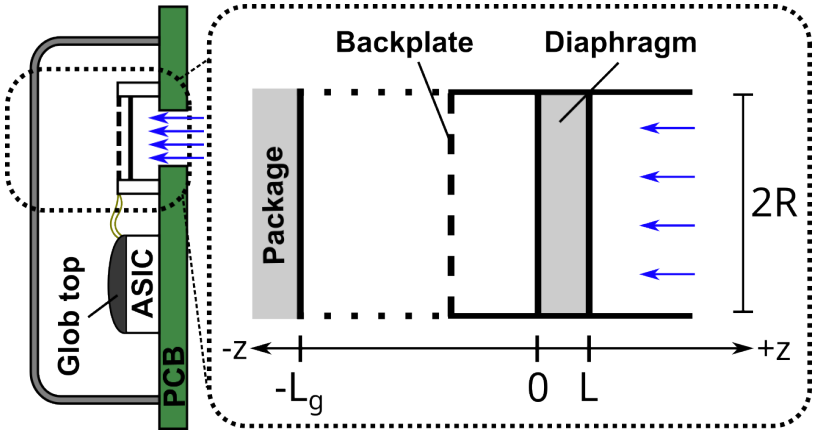

Figure 4.6: The coordinate system used for the MEMS microphone model.

smart home devices were vulnerable to optical powers much lower than our 60mW high power experiment, so successful injection was possible at the maximum range of 50 meters. Notice that the final four devices are smartphones and tablets, with limited voice authentication capabilities. These devices were much less vulnerable than the other devices, and could only be affected at less than 20 meters.

Table 4.1: Tested devices with minimum successful power and maximum distance achievable at the given power of 5 mW and 60 mW. A 110 m long hallway was used for 5 mW tests while a 50 m long hallway was used for tests at 60 mW.

Device	Backend	Category	Authentication	Minimum Power [mW]	Max Distance at 60 mW [m]*	Max Distance at 5 mW [m]**
Google Home	Google Assistant	Speaker	No	0.5	50+	110+
Google Home Mini	Google Assistant	Speaker	No	16	20	—
Google Nest Cam IQ	Google Assistant	Camera	No	9	50+	—
Echo Plus 1st Generation	Alexa	Speaker	No	2.4	50+	110+
Echo Plus 2nd Generation	Alexa	Speaker	No	2.9	50+	50
Echo	Alexa	Speaker	No	25	50+	—
Echo Dot 2nd Generation	Alexa	Speaker	No	7	50+	—
Echo Dot 3rd Generation	Alexa	Speaker	No	9	50+	—
Echo Show 5	Alexa	Speaker	No	17	50+	—
Echo Spot	Alexa	Speaker	No	29	50+	—
Facebook Portal Mini (Front Mic)	Alexa	Speaker	No	1	50+	40
Facebook Portal Mini (Front Mic)***	Portal	Speaker	No	6	40	—
Fire Cube TV	Alexa	Streamer	No	13	20	—
EcoBee 4	Alexa	Thermostat	No	1.7	50+	70
iPhone XR (Front Mic)	Siri	Phone	Yes	21	10	—
iPad 6th Gen	Siri	Tablet	Yes	27	20	—
Samsung Galaxy S9 (Bottom Mic)	Google Assistant	Phone	Yes	60	5	—
Google Pixel 2 (Bottom Mic)	Google Assistant	Phone	Yes	46	5	—

*Data limited to a 50 m long corridor, **Data limited to a 110 m long corridor, ***Data generated using only the first 3 commands.

4.3.3 A Model of LSI in MEMS microphones

While *Light Commands* showed that LSI on MEMS microphones was possible, the underlying mechanisms causing the effect were mysterious. In order to explain LSI and potentially defend against future attacks, a model of the injection is presented. The effects of laser signal injection into MEMS microphones can be described by the combination of three different physical mechanisms:

- Thermoelastic (TE) Bending [35, 25]: As the MEMS structures absorb the incoming light, they heat up and expand. This displaces the diaphragm when a bending moment is generated by thermal asymmetries within the diaphragm.
- Thermal Diffusion (TD) [38, 34]: As the diaphragm heats up, it also heats the surrounding air. The periodically heated air column expands adiabatically, generating a pressure wave that displaces the diaphragm.
- Photovoltaic Effect (PV) [45]: As light interacts with semiconductor components, it generates excess charge carriers. When these charge carriers appear

within p-n junctions on the ASIC, it will generate a photocurrent and voltage signal that will be coupled to the output voltage of the microphone.

While several other phenomena described in Section 2.3 can potentially contribute to a signal on the output of the microphone , we did not see a significant contribution outside of these three mechanisms. A full description of these extra effects is given in Appendix A The models of these mechanisms rely on a coordinate system defined in Figure 4.6. These effects are summarized in Figure 4.7. Note that all physical parameters within these models are described by the real parts of any complex expressions.

4.3.3.1 Optical Irradiance Model

In the LSI attack, the optical power P of the attacking laser was modulated to inject an audio signal. By modulating the light, the result was a change in light irradiance (optical power density) entering the acoustic port of the microphone. Without any changes to aiming or focus, a single frequency component ω of the irradiance signal entering the acoustic port can be modeled as the real part of the complex expression:

$$I = I_B + I_0 e^{j\omega t} = [P_B + P_0 e^{j\omega t}]/A_B \tag{4.2}$$

where P_0 is the amplitude of the sinusoidal power signal entering the acoustic port, P_B is the bias on the optical power signal, and A_B is the cross-sectional area of the laser beam that enters the acoustic port (A_B will be equal to the cross-sectional area of the acoustic port in most attack cases). To ensure the model is linear , P_0 must be less than P_B. Besides this condition, however, the contribution of the bias signal I_B can often be ignored. Therefore, for the rest of this section, we will only consider the time-varying portion of this signal: $I_0 e^{j\omega t}$.

Assuming the angle of incidence is normal to the diaphragm, the MEMS diaphragm will reflect (I_R), transmit (I_T), and absorb (I_A) a certain amount of incoming

irradiance depending on the wavelength λ of the incoming light:

$$I_R = R_\lambda I_0 e^{j\omega t} \tag{4.3}$$

$$I_T = T_\lambda I_0 e^{j\omega t} \tag{4.4}$$

$$I_A = \left[1 - R_\lambda - T_\lambda\right] I_0 e^{j\omega t} \tag{4.5}$$

where (R_λ) and (T_λ) are the optical reflectance and transmittance dependent on the light wavelength λ.

The amount of light that is reflected, transmitted, and absorbed is highly dependent on the materials and structure of the MEMS device and can quickly become difficult to model at a high optical transmittance (T_λ). In general, for the materials used in MEMS structures (such as polysilicon or aluminum nitride), the transmittance monotonically decreases as the wavelength of the incoming light decreases. This means that shorter wavelength "bluer" light will be absorbed and blocked much more than longer wavelength "redder" light. This is vitally important for understanding the contributions towards each of the physical effects on the microphones.

4.3.3.2 Thermoelastic Effects (TE)

The thermoelastic bending component of the photoacoustic signal results from thermal moments that are generated as the diaphragm is heated by the incoming laser signal. This effect was first modeled by Rousset et al. [35] from thermal moments arising from an asymmetric heat distribution through the thickness of a heated plate. In the case of MEMS diaphragms, however, the structures are thin and insulated by air, which has a relatively low thermal conductivity. This means that the heat diffuses through the thickness almost immediately, causing the diaphragm to be at a nearly

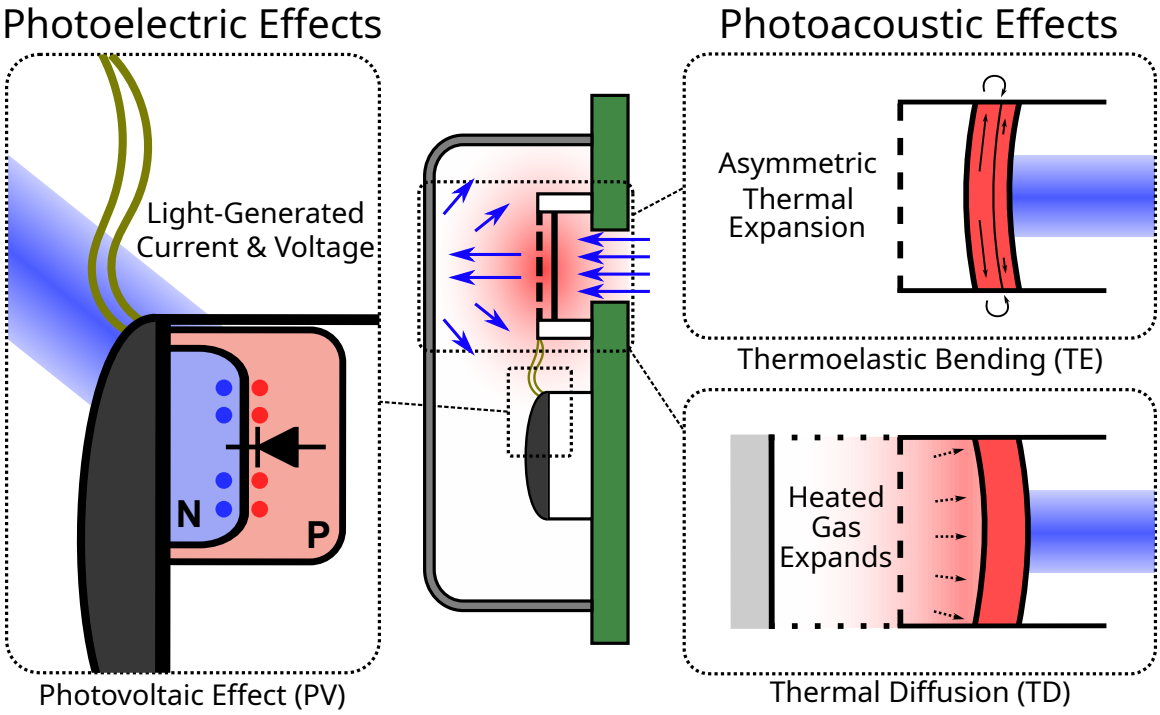

Figure 4.7: A summary of the three primary physical phenomena that were investigated in this work. Two mechanisms are photoacoustic and dependent on the heating of the diaphragm and the air. The last one is photoelectric and dependent on carrier generation within the ASIC.

uniform temperature. The temperature of the diaphragm (T_d) can be modeled by:

$$T_d \approx \frac{I_A}{j\omega\rho c_p L} \tag{4.6}$$

where ρ is the density of the diaphragm material, c_p is the specific heat capacity, and L is the thickness of the diaphragm. Note that the temperature decreases as the frequency of modulation ω increases.

While Rousset's model does not predict bending when there is a uniform temperature, there are still cases where bending can occur. Rather than an asymmetry of temperature, Todorović et al. [25] described the effects of an asymmetry of material properties between a substrate and a thin film. The differences in elasticity shift the neutral plane (z_n), and the differences in thermal expansion generate a moment

even in the case of uniform temperature. Using this model, a quasistatic analysis can represent the average displacement of the membrane with:

$$w_{TE} = \frac{1}{4} R^2 M_T T_d \tag{4.7}$$

where R is the radius of the diaphragm, and M_T is a moment-generating constant determined from the thermal properties of the materials and geometry of the diaphragm:

$$M_T = \frac{\int_0^L \alpha_T(z) E(z - z_n) dz}{\int_0^L E(z)(z - z_n)^2 \, dz} \tag{4.8}$$

$$z_n = \frac{\int_0^L E(z) \; z \; dz}{\int_0^L E(z) \; dz} \tag{4.9}$$

where $\alpha_T(z)$ and $E(z)$ are the linear thermal expansion coefficient and Young's modulus of the material at each coordinate through the diaphragm thickness L.

The important factors to consider for TE are that it is dependent only on the static material properties of the diaphragm, the frequency of the incoming modulation, and the light that is absorbed by the diaphragm. Shorter wavelengths will be absorbed more strongly by the diaphragm, but any absorbed light will generate a bending effect. Since the diaphragm temperature T_d is inversely proportional to the modulation frequency, the TE component will decrease by a factor of ω^{-1} as the modulation frequency increases.

4.3.3.3 Thermal Diffusion (TD)

Thermal diffusion or the "thermal-piston" effect is the process by which an acoustic pressure wave is generated as the air in the microphone periodically heats and expands. This effect was first described by Rosencwaig and Gersho [38] while performing photoacoustic experiments on a closed cell of air. In their model, an incoming

laser signal is absorbed by a surface, causing a sudden increase in the temperature of the surface. This heat then diffuses into the surrounding air. Because the air has a much lower thermal diffusivity than the absorbing surface, the heat conducts slowly, resulting in a layer of hot air close to the surface. This layer of air will expand adiabatically pushing out against the rest of the air within the closed cell and generating a pressure signal.

In the context of MEMS microphones, a pressure signal will be generated within the package and push the diaphragm outward. To determine the displacement of the microphone, we start by first modeling the average temperature of the gas column that extends from the MEMS structure to the back package. The average temperature of the air column can be calculated as described in [34] using a one-dimensional heat transfer model, and using the temperature of the diaphragm and the ambient temperature as the boundary conditions. The spatially-averaged gas temperature reduces to:

$$T_g = \frac{\tanh\left(L_g\sigma_g/2\right)}{L_g\sigma_g}T_d \tag{4.10}$$

$$\sigma_g = \sqrt{j\omega/\alpha_g} \tag{4.11}$$

where L_g is the height of the gas column (the distance from the MEMS structure to the package), and σ_g is the complex thermal diffusion parameter defined from the modulation frequency ω and the thermal diffusivity of the air α_g.

Using the average temperature of the air column, the pressure signal can then be calculated from the ideal gas law with adiabatic expansion:

$$P = \frac{\gamma P_0 V_g}{V_0 T_0}T_g \tag{4.12}$$

$$V_g = \pi R^2 L_g \tag{4.13}$$

where P_0, V_0, and T_0 are the ambient pressure, volume of the microphone's back cavity, and ambient temperature respectively. The constant γ is the ratio of specific heats of air at constant pressure and constant volume, which can be approximated as 7/5. The volume of the heated air V_g is defined from the area of the MEMS structure and the height of the heated air column L_g.

From this pressure, the quasi-static average displacement of the diaphragm at low frequencies can be described with:

$$w_{TD} = \frac{A_r \pi R^2}{K_d} P \tag{4.14}$$

where K_d is the effective spring constant of the diaphragm and A_r is the effective acoustic area coefficient, which accounts for the differences in displacement and pressure at each point on the diaphragm. The coefficient A_r is a number from 0 to 1, and it is dependent on the mechanical boundary conditions of the diaphragm. The mechanical model of the diaphragm has been explored in [111] and [112].

The TD component is primarily dependent on the temperature of the diaphragm and the acoustic properties of the microphone system. Notice that for low modulation frequencies, the heat diffuses entirely through the gas column, and the average temperature can be approximated as $T_g = T_d/2$, and therefore the signal decreases at a rate of ω^{-1}. As the modulation frequency increases, however, the heat only diffuses partway through the column, reducing the average temperature to $T_g = T_d/l_g\sigma_g$. In this case, the TD signal is proportional to $\omega^{-3/2}$, as σ_g is proportional to $\omega^{1/2}$. Therefore the TD signal gets much weaker than other effects as the modulation frequency increases. Beyond the modulation frequency, the output signal will be directly proportional to the ambient pressure and inversely proportional to the ambient temperature. This can be used to isolate the TD phenomenon, as it is the only one to be primarily affected by changes to these ambient conditions.

4.3.3.4 Photovoltaic Effects (PV)

The laser will affect the output voltage of the microphone by inducing a photocurrent within the signal measurement and processing circuits on the ASIC. The photocurrent generated by the incoming laser signal can be represented by [45]:

$$I_\phi \approx G_R \eta_\lambda \frac{\lambda I_T}{hc} \tag{4.15}$$

where G_R is a gain factor dependent on the optical, material, and electrical properties of the device, η_λ is the quantum efficiency dependent on the light wavelength λ, h is the Planck constant, and c is the speed of light. While many device-specific factors influence G_R and make it difficult to calculate, the quantum efficiency can be approximated as unity except at wavelengths with insufficient energy to generate charge carriers within the ASIC materials. For silicon, light wavelengths longer than approximately 1100 nm are unable to excite electrons above the 1.12 eV band gap, meaning that no charge carriers can be generated.

From this model, we can see that the important factor affecting the photovoltaic signal is the wavelength of the incident light. Longer light wavelengths will increase the generated photocurrent both by transmitting more light through the diaphragm, and by having more carrier-generating photons per unit of power. This trend continues until the photons do not contain enough energy to excite charges passed the band gap (e.g. 1100nm for silicon). At this point, the quantum efficiency and photocurrent will drop to zero.

4.3.3.5 A Model of the Microphone Output

The output voltage V_{out} of the MEMS microphone can be described as a linear combination of the displacement of the sensing diaphragm (w) and the photocurrent generated as light interacts with the sensitive components on the ASIC (I_ϕ). This

can be represented as:

$$V_{out} \approx G_w w + G_\phi I_\phi \tag{4.16}$$

where G_w and G_ϕ are the gain factors that translate the displacement and photocurrent into voltage respectively. These gain factors are dependent on many different properties in each MEMS microphone and are difficult to calculate and measure without specific knowledge of the ASIC design. Ultimately, their exact values are unnecessary in determining the cause of the output signal that we see.

The displacement of the diaphragm w can be modeled as a linear combination of the two photoacoustic effects:

$$w = w_{TE} + w_{TD} \tag{4.17}$$

where w_{TE} and w_{TD} are the displacement due to thermoelastic bending, plasmaelastic bending, and thermal diffusion respectively.

Because the output voltage is a combination of these different factors, it can be difficult to isolate and describe any one factor and its contribution to the output signal. That is why we present a set of experiments to isolate and measure the contribution of each effect in the next section.

4.3.4 A Setup to Investigate LSI in MEMS Microphones

Now that the relevant physical effects have been modeled, it is seen that each of the effects can be separated by specific parameters such as air pressure, light wavelength, and frequency. Now experiments can be made to determine the contributions of each effect on a set of eight microphones. To perform this investigation, a setup was constructed with a specific set of capabilities to isolate the potential physical phenomena affecting the MEMS microphones.

Table 4.2: The MEMS microphones used in experiments

Device	Manufacturer	Type	Output	Diaphragm	Globtop on ASIC
CMM3526	CUI Devices	Capacitive	Analog	Front	✓
SPU0410	Knowles	Capacitive	Analog	Front	✓
ICS41350	InvenSense	Capacitive	Digital	Back	✓
ADMP401	Analog Devices	Capacitive	Analog	Back	-
SPA1687	Knowles	Dual Capacitive	Analog	Front	✓
SPH0641	Knowles	Dual Capacitive	Digital	Front	✓
VM1010	Vesper	Piezoelectric	Analog	Single	✓
VM3000	Vesper	Piezoelectric	Digital	Single	✓

4.3.4.1 Target Microphones

Eight different MEMS microphones were selected as targets for our experiments. These microphones are summarized in Table 4.2. Out of the eight targets, six of them are capacitive-sensing, while the Vesper VM1010 and VM3000 are piezoelectric-sensing microphones. Three of the microphones (ICS41350, SPH0641, and VM3000) have digital Pulse Density Modulation (PDM) outputs, demonstrating that the laser injection affects the devices even when there is a digital output. All of the capacitive-sensing microphones have doped polysilicon diaphragms [113, 114], while the Vesper microphones have diaphragms consisting primarily of aluminum nitride [115]. The SPH0641 and the SPA1687 each have two diaphragms-backplate pairs instead of a single pair. All of the microphones except the ADMP are roughly the same package size with the same volume of air in the back cavity. The ADMP401 was included to show how a microphone is affected when there isn't any light-blocking globtop, as it was the only microphone we targeted that did not have any light protection. The ADMP also contained external amplification and filtering circuitry, but the displayed results are the signal from the microphone directly.

4.3.4.2 Signal Conditioning and Measurement

The voltage output from each of the target microphones was measured with a Stanford Research Model SR560 pre-amplifier connected to a Picoscope 5444D oscilloscope. The preamplifier was set to a low-pass filter with a 30kHz cutoff frequency and a -6dB/octave roll-off. All the microphones were powered with a Sigilent SPD3303C power supply set to a constant +3V. In the case of the digital output microphones, a 0-3V 2.4MHz clock signal was generated with a Tektronix AFG3102 function generator. To convert the digital PDM signal to an analog waveform, a simple RC low-pass filter consisting of a 1 kΩ resistor and a 4.7 nF capacitor was attached to the output. A Dell XPS laptop running a custom MATLAB program with the Instrument Control Toolbox was used to obtain data from the Picoscope while simultaneously controlling the laser output with a connection to the function generator.

4.3.4.3 Controlling Optical Irradiance

In order to have precise control for our experiments, we used several tools to control the optical power and focus of the laser output. We used five laser diodes in our experiments: a 1470 nm Mitsubishi ML920J16S, a 904 nm Thorlabs L904P010, a 638 nm Thorlabs L638P150, and a 450 nm Osram PLT5 450B. Since the optical output power of a laser diode is linearly proportional to the current across the junction, the optical power can be controlled with a variable current source. This current source was formed with a laser driver connected to a function generator. In our setup, we used a Thorlabs LDC205C laser driver, controlled by a Tektronix AFG3102 function generator. During experimentation, a Dell XPS laptop running a custom MATLAB program with the Instrument Control Toolbox was used to generate a frequency sweep on the function generator while simultaneously capturing data. A Thorlabs PM100USB power meter with an S425C head was used to measure and calibrate the optical power output of the 1470 nm laser. An S121C head was used to calibrate the

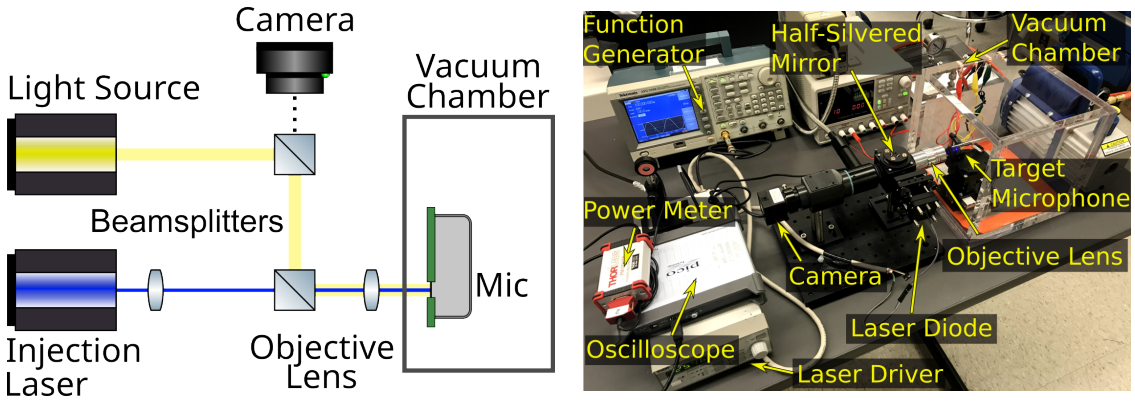

Figure 4.8: A setup to precisely measure LSI in MEMS microphones

optical power of the rest of the laser diodes.

We developed a setup to control the aiming and focus of the laser beam using a C-mount camera, a Thorlabs LDH56-P2 laser collimation cage, two half-silvered mirrors as beamsplitters, and a Mitutoyo 5x objective lens. A Hayashi LA-100USW was used as a light source to assist in viewing the target diaphragm, but it was powered down during experimentation. In the case of the 1470 nm laser, a Thorlabs VRC2 detector card was required to aim and focus the beam. The full optical setup is shown in Figure 4.8. This setup allowed us to visually see the focus and position of the laser beam as it was injected into the MEMS acoustic port.

4.3.4.4 Vacuum Setup

In order to test the effects of low atmospheric pressure, we performed a signal injection attack while the microphone was in a vacuum chamber. Figure 4.8 shows an overview of the setup. We used a BVV vacuum chamber with acrylic transparent walls and an included pressure gauge. A Zeny VP125 vacuum pump was used to evacuate air from the chamber. A Thorlabs 3-Axis manual stage with rotation was used to hold the microphone, allowing for fine control of the position and rotation of the target before turning on the vacuum. Strips of thin copper tape were used to transport signals in and out of the vacuum chamber.

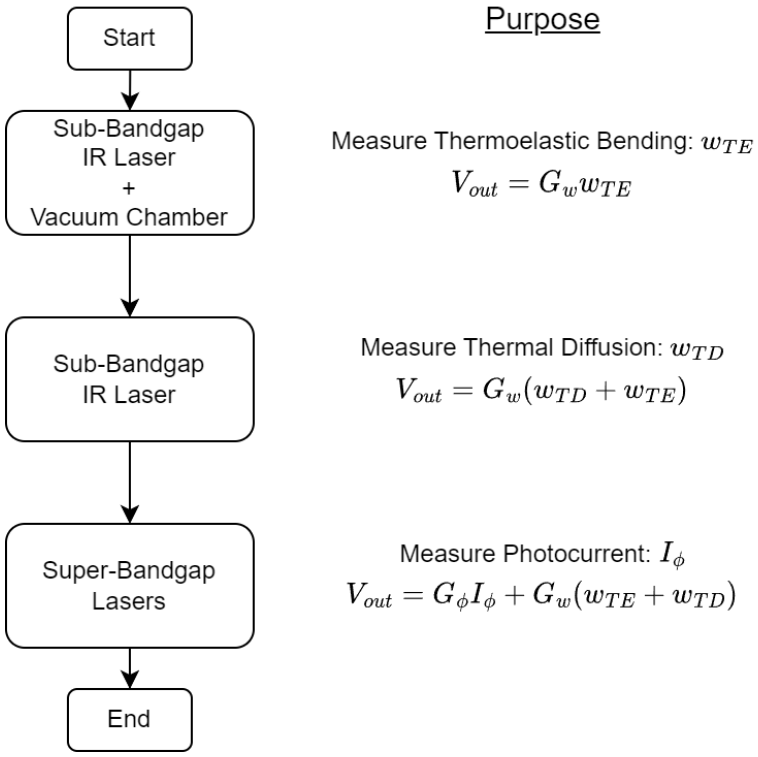

Figure 4.9: The experimental procedure to determine the contributions of the three physical mechanisms to the output voltage of each microphone.

4.3.5 Characterizing LSI in Commercial Microphones

With a flexible and precise setup to isolate various physical phenomena, a procedure was developed to identify the contributions of each of the applicable physical phenomena on the output signal of a microphone under a laser signal injection attack. In our experiments, we found that the different designs of each microphone lead to different dominant factors being exploited in a laser signal injection attack. The purpose of this set of experiments is to measure the relative contribution of each physical phenomenon towards the output signal. With an understanding of these phenomena, it will be possible to make design changes to reduce the contribution of each effect, leading to MEMS microphones that are more resistant to laser signal injection. To demonstrate this procedure, we perform the characterization on eight different commercial MEMS microphones with different vendors and properties.

An overview of this procedure is shown in Figure 4.9. The main idea is to isolate a single physical effect and measure the amplitude and phase response of the output signal as we sweep the modulation frequency of the input laser. We do this because the amplitude and phase responses form a way to identify each component, as each one has a different response to the incoming light signal. Once we have a way to identify a single component, we add each of the other two components to determine their relative contribution to the output signal. By comparing the shifts in amplitude and phase, it identifies when each component is dominant so we can obtain a characterization of the laser signal injection.

Eight different MEMS microphones were selected as targets for our experiments. These microphones are summarized in Table 4.2. Out of the eight targets, six of them are capacitive-sensing, while the Vesper VM1010 and VM3000 are piezoelectric-sensing microphones. Three of the microphones (ICS41350, SPH0641, and VM3000) have digital Pulse Density Modulation (PDM) outputs, demonstrating that the laser injection affects the devices even when there is a digital output. All of the capacitive-sensing microphones have doped polysilicon diaphragms [113, 114], while the Vesper microphones have diaphragms consisting primarily of aluminum nitride [115]. The SPH0641 and the SPA1687 each have two diaphragms-backplate pairs instead of a single pair. All of the microphones except the ADMP are roughly the same package size with the same volume of air in the back cavity. The ADMP401 was included to show how a microphone is affected when there isn't any light-blocking globtop, as it was the only microphone we targeted that did not have any light protection. The ADMP also contained external amplification and filtering circuitry, but the displayed results are the signal from the microphone directly.

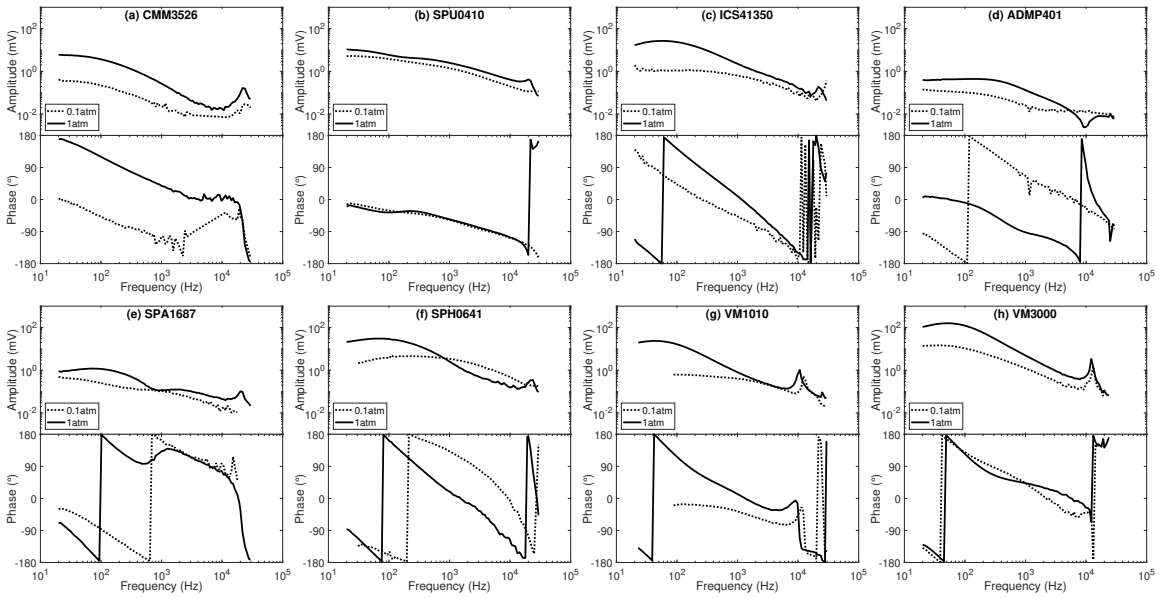

Figure 4.10: The results from the vacuum chamber experiments with a sub-bandgap IR laser. All measurements were completed

4.3.5.1 Determining the Contribution of Thermal Effects with a Sub-Bandgap IR Laser and a Vacuum Chamber

The first step in the process to characterize the laser signal injection is to isolate a single physical phenomenon that could be generating the effects that we see on the output of the microphone. This effect was chosen to be the thermoelastic bending effect. This leaves the two other effects to be removed: photovoltaic effects and thermal diffusion.

To remove the photovoltaic signal, we performed a laser signal injection experiment using a 1470 nm IR laser, which has much less photon energy than the band gap of silicon. These experiments are inspired by thermal laser stimulation (TLS) failure analysis [44], where laser wavelengths longer than 1100 nm are used to generate heat but do not have enough photon energy to generate a photocurrent (See Section 4.3.3.4). By using a 1470 nm laser for signal injection, we ensure that thermal effects are the only phenomena generating the output signal of the microphone.

While the silicon diaphragms will be fairly transparent to this wavelength of light, we found that enough energy is absorbed to have a measurable signal.

Now that we have isolated down to the thermal effects, we can differentiate between thermoelastic bending and thermal diffusion using a vacuum chamber. As mentioned in Section 4.3.3.3, the thermal diffusion effect is directly proportional to the ambient pressure of the air within the microphone. When we use a vacuum chamber to reduce the ambient pressure, we also reduce the contribution of the thermal diffusion effect, as there is less air to generate photoacoustic waves. When the vacuum chamber is combined with the 1470nm laser, we can isolate the effects of thermoelastic bending alone. Once we have the thermoelastic bending component, we can determine the thermal diffusion component by reintroducing air into the vacuum chamber while still performing the signal injection. The difference between the signal at low pressure, and the signal at atmospheric pressure will give us the thermal diffusion component.

We performed these laser signal injection experiments on eight different commercial microphones described in Section 4.3.4.1 and Table 4.2. The laser was kept at a bias power of 5 mW and an injection signal amplitude of 1 mW. A frequency sweep of the injection laser was used to collect the amplitude and phase response of the output voltage signal. The changes in amplitude and phase of provide a way to identify when each thermal component is dominant. We performed the experiment first at a pressure of 0.1 atm, and then repeated the experiment at 1 atm.

The results of these experiments are presented in Figure 4.10. The eight microphones are presented in eight separate subfigures. The top of each subfigure shows the amplitude response of the microphone output over the selected frequencies, while the bottom of each subfigure shows the phase response. Notice that each of the microphones had a measurable thermoelastic bending signal while under vacuum, indicating that all of the diaphragms had some innate thermal asymmetry that was

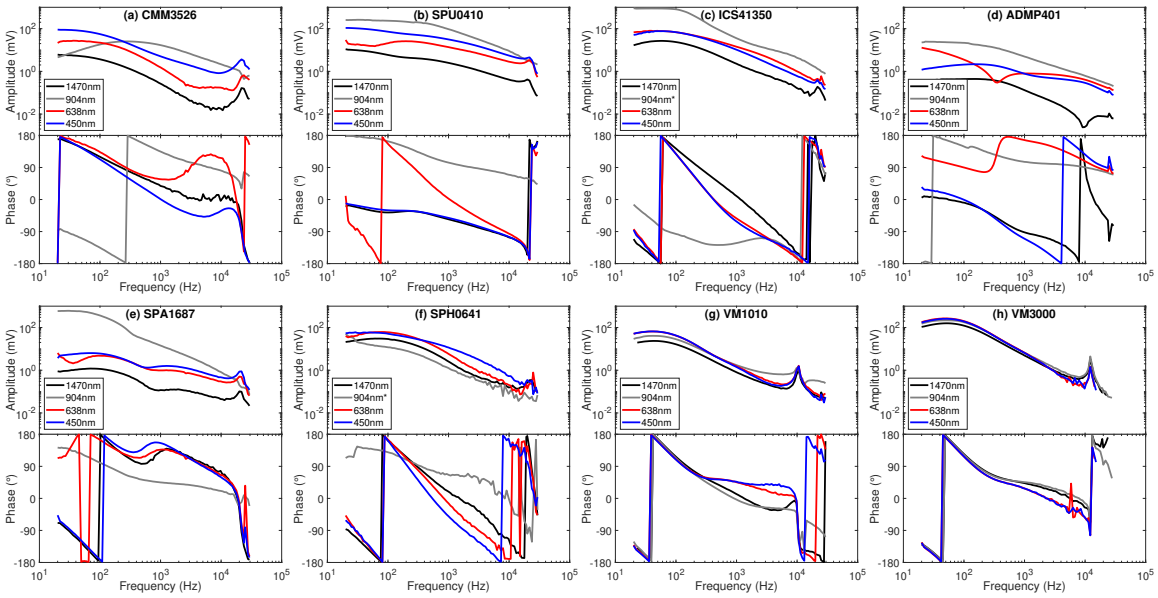

Figure 4.11: A comparison of photoacoustic and photoelectric effects on MEMS microphones at 5 mW bias and a 1mW amplitude laser signal. The effects of the 1470nm laser are entirely due to thermal effects, while the rest will be some mixture of thermal and electric effects. (*)Asterisks indicate that the injected power was at 0.2 mW bias and 0.1 mW amplitude to prevent disabling of the microphone.

being exploited by the laser signal injection. In fact, the output of the Knowles SPU0410 (Fig. 4.10b) seems to be driven nearly entirely by thermoelastic bending. The rest of the microphones had significant changes to their output voltage signals as we introduced air back into the chamber, indicating the presence of thermal diffusion effects. For all of the microphones except the SPU0410, the TD component is clearly dominant over the TE component. This is demonstrated by a significant increase in output amplitude, as well as a change in the output phase, especially at low frequencies. At these low frequencies, the thermal diffusion component is out-of-phase with the thermoelastic bending component. As frequencies increase, the thermal diffusion signal gets weaker, and the output signal aligns with the thermoelastic bending signal.

4.3.5.2 A Comparison of Photoacoustic Effects and Photoelectric Effects

Now that we have a clear idea of the contribution of photoacoustic effects in each of the microphones, the next step in the process is to determine the contribution of photocurrent generation. To do this, we repeat the modulation frequency sweep experiment in Section 4.3.5.1 using three more wavelengths of an injection laser: a 904nm laser, a 638nm laser, and a 450nm laser. All three of these wavelengths have energies above the band gap of silicon and will produce a photocurrent on sensitive parts of the ASIC circuitry. Because the shape of the amplitude and phase responses of the thermal effects will stay the same, any changes to the amplitude and phase response can be attributed to photocurrent generation.

The results of our experiments are shown in Figure 4.11. We again perform the experiments on the same eight microphones. Most of these experiments were performed with a bias power of 5mW with a 1mW amplitude signal. Here we compare all three laser wavelengths above the bandgap with the sub-bandgap laser. As we can see, nearly all microphones exhibit some photoelectric effects in the ASIC. This is especially apparent when using the 904nm laser due to its long wavelength and high diaphragm transmission, which contributes to the highest PV signal as described in Section 4.3.3.4. For most of the microphones, photocurrent generation dominates photoacoustic effects for 904nm light. For the 450nm laser, the trend is reversed, and the results follow very closely with the 1470nm sub-bandgap laser. This indicates that for blue light, the signal is entirely driven by photoacoustic effects. The red 638nm laser shows how the photoacoustic and photoelectric signals mix in many of the microphones.

To understand the characteristics of the photocurrent generation, the primary result to consider is the experiments with the 904nm laser. Notice that the output amplitude of the PV signal is considerably higher than the photoacoustic signal for most of the microphones. Just like the photoacoustic signal, the PV signal exhibits a

decrease in signal amplitude at higher frequencies. This is likely due to the electrical properties of the photosensitive parts of the ASIC, which are difficult to predict. From cases where the PV signal is dominant, we can see that the PV signal tends to be out-of-phase with the photoacoustic signal, leading to clear antiresonances in some of the microphones where the photoelectric and photoacoustic effects are competing for dominance. This is especially apparent for the ADMP401 (Fig. 4.11d), where the output from the red 638 nm laser injection has an antiresonance at 300 Hz where the PV signal is completely out-of-phase with the photoacoustic signal.

For two of the microphones, the ICS41350 (Fig. 4.11c) and the SPH0641 (Fig. 4.11f), we performed our experiments with the 904nm laser at a reduced bias power of 0.2mW and an amplitude of 0.1mW. We did this because the generated photocurrent was so significant that it could actually disable the output of the microphone at high enough injection power. While the exact explanation of why this occurs is difficult without a full understanding of the ASIC circuitry, we believe that it is due to the photocurrent causing a short circuit in a vital signal processing component. After the laser is turned off, the device quickly returns to normal operation. This effect of temporarily disabling the microphone output with a laser has significant consequences on the security of systems using them.

4.3.5.3 The Causality of LSI in Commercial Microphones

While we encourage microphone designers to determine vulnerabilities within each of their devices, we wanted to discuss the general vulnerability trends that we found during our research on these eight commercial microphones. A summary of our results is shown in Table 4.3.5.3. For each wavelength of incoming light, we ranked each measurable contribution of each effect from 1 (the dominant contribution) to 3 (the least contribution). This was repeated for each microphone to show how the different effects combine to produce the output voltage signal.

Table 4.3: A ranking of the contribution of each physical effect on the output amplitude (1=strongest, 3=weakest, an asterisk (*) denotes the microphone is temporarily disabled).

	IR 1470nm			IR 904nm			Red 638nm			Blue 450nm		
Device	TE	TD	PV	TE	TD	PV	TE	TD	PV	TE	TD	PV
CMM3526	-	1	-	-	2	1	-	1	2	-	1	-
SPU0410	1	2	-	-	-	1	1	3	2	1	3	2
ICS41350	2	1	-	2	-	1*	2	1	-	2	1	-
ADMP401	-	1	-	-	-	1	-	2	1	-	1	2
SPA1687	2	1	-	-	-	1	2	1	3	2	1	-
SPH0641	-	1	-	-	-	1*	-	1	-	-	1	-
VM1010	-	1	-	-	1	2	-	1	-	-	1	-
VM3000	-	1	-	-	1	-	-	1	-	-	1	-

The primary concern in most of the microphones is the photocurrent generated on the ASIC when using super-bandgap IR wavelength lasers. These longer wavelengths penetrate the polysilicon MEMS structures to affect the photosensitive ASIC directly. While the ASIC of almost every microphone we examined contained an opaque globtop covering for environmental protection, the globtop did not adequately protect against IR light. The two piezoelectric-sensing microphones we investigated were much more resistant to photocurrent generation, as the IR light was likely unable to penetrate the aluminum nitride diaphragm.

The secondary concern in the microphones we investigated was the thermal diffusion photoacoustic signal. While biased towards low frequencies, the thermal diffusion signal was significantly stronger than the other photoacoustic effects in all but one of the microphones. This is especially apparent while using visible light lasers, where the majority of the incoming light is absorbed by the MEMS diaphragm. This is also the primary concern for the Vesper piezoelectric microphones, as their diaphragms absorb nearly all incoming laser signals.

Finally, several microphones exhibit a strong thermoelastic bending signal within certain frequency regions of the injected signal. This is especially apparent in the

SPU0410, where the thermoelastic bending signal is the dominant photoacoustic effect at all frequency regions. For all of the other microphones, thermoelastic bending is present but often overridden by the thermal diffusion signal, especially for low frequencies.

4.4 Consequences on Voice-Controllable Systems

The results of the previous sections clearly demonstrate the feasibility of LSI of voice commands into voice-controlled devices across large attack distances. In this section, we explore the security implications of such an injection, as well as experiment with more realistic attack conditions.

4.4.1 A Low-Power Cross-Building Attack

For the long-range attacks presented in Section 4.3.2, we deliberately placed the target device so that the microphone ports are facing directly into the laser beam. While this is realistic for some devices that have microphone ports on their sides, such an arrangement is artificial for devices with top-facing microphones (unless mounted sideways on the wall).

In this section, we perform the attack under more realistic conditions, where an attacker aims from another higher building at a target device placed upright on a window sill.

We used the laser diode, telephoto lens, and laser driver from Section 4.3.2, operating the diode at 5 mW (equivalent to a laser pointer) with the same modulation parameters as in the previous section. Next, we placed a Google Home device (which only has top-facing microphones) upright near a window, in a fourth-floor office (15 meters above the ground). The attacker's laser was placed on a platform inside a nearby bell tower, located 43 meters above ground level. Overall, the distance between the attacker and laser was 75 meters, see Figure 4.12 for the configuration.

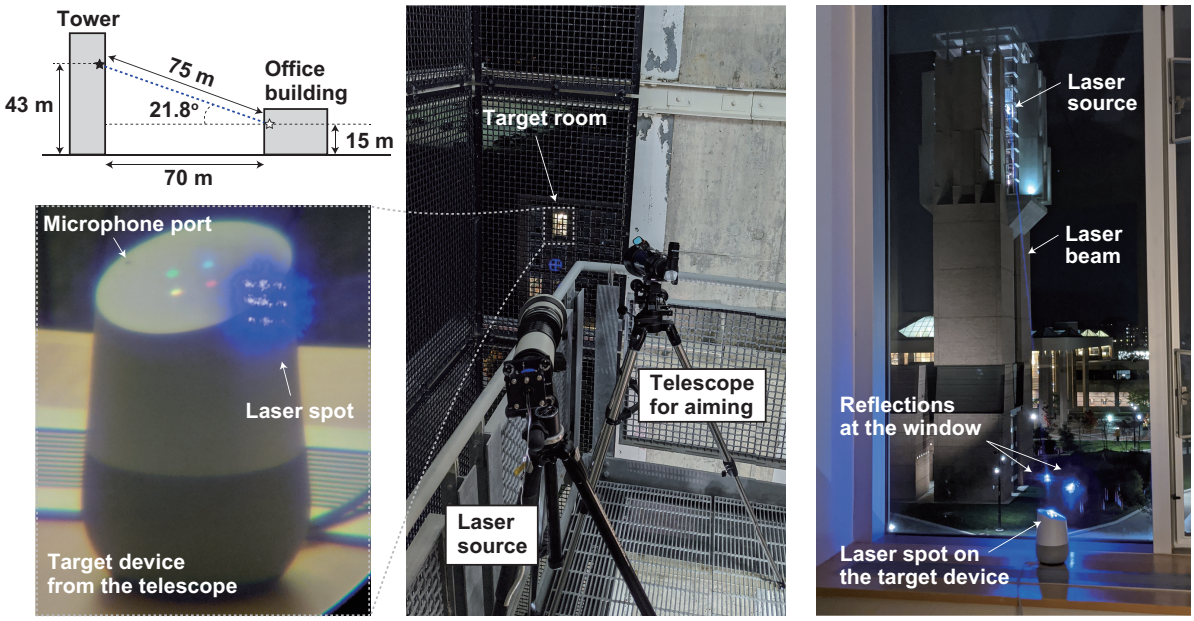

Figure 4.12: Setup for the low-power cross-building attack: (Top left) Laser and target arrangement. (Bottom left) Picture of the target device as visible through the telescope, with the microphone ports and laser spot clearly visible. (Middle) Picture from the tower: laser on telephoto lens aiming down to the target. (Right) Picture from the office building: laser spot on the target device.

As in Section 4.3.2, it is impossible to focus the laser using the small lens typically used for laser pointers. We thus mounted the laser to an Opteka 650-1300 mm telephoto lens. Next, to aim the laser across large distances, we mounted the telephoto lens on a Manfrotto 410 geared tripod head. This allows us to precisely aim the laser beam on the target device across large distances, achieving an accuracy far exceeding the one possible with regular (non-geared) tripod heads where the attacker's arm directly moves the laser module. Finally, in order to see the laser spot and the device's microphone ports from far away, we have used a consumer-grade Meade Infinity 102 telescope. As can be seen in Figure 4.12 (left), the Google Home microphone's ports are clearly visible through the telescope.

We have successfully injected commands into the Google Home target in the above-described conditions. We note that despite its low 5 mW power and windy condi-

tions (which caused some beam wobbling due to laser movement), the laser beam successfully injected the voice command while penetrating a closed double-pane glass window. While causing negligible reflections, the double-pane window did not cause any visible distortion in the injected signal, with the laser beam hitting the target's top microphones at an angle of 21.8 degrees and successfully injecting the command without the need for any device- or window-specific calibration. We thus conclude that cross-building laser command injection is possible, at large distances and under realistic attack conditions.

4.4.2 Exploring Stealthy Attacks

The attacks described so far can be spotted by the user of the targeted VCS in three ways. First, the user might notice the light indicators on the target device following a successful command injection. Next, the user might hear the device acknowledging the injected command. Finally, the user might notice the spot while the attacker tries to aim the laser at the target microphone port.

While the first issue is a limitation of our attack (and in fact of any command injection attack), in this section we explore the attacker's options for addressing the remaining two issues.

To tackle the issue of the device owner hearing the targeted device acknowledging the execution of voice command (or asking for a PIN number during the brute forcing process), the attacker can start the attack by asking the device to lower its speaker volume. For some devices (EcoBee, Google Nest Camera IQ, and Fire TV), the volume can be reduced to completely zero, while for other devices it can be set to barely-audible levels. Moreover, the attacker can also abuse device features to achieve the same goal. For Google Assistant, enabling the "do not disturb mode" mutes reminders, broadcast messages, and other spoken notifications. For Amazon Echo devices, enabling "whisper mode" significantly reduces the volume of the device

responses during the attack to almost inaudible levels.

The attacker can also use an invisible laser wavelength to avoid having the owner spot the laser light aimed at the target device. However, as the laser spot is also invisible to the attacker, a camera sensitive to the appropriate wavelength is required for aiming. Experimentally verifying this, we replicated the attack on a Google Home device using a 980-nm Infrared (IR) laser (Lilly Electronics 30 mW laser module). We then connected the laser to a Thorlabs LDC205C driver, limiting its power to 5 mW. Finally, as the spot created by IR lasers is invisible to humans, we aimed the laser using a smartphone camera (as these typically do not contain infrared filters).

Using this setup, we have successfully injected voice commands to a Google Home at a distance of about 30 centimeters in the same setup as Section 4.3.2. The spot created by the infrared laser was barely visible using the phone camera, and completely invisible to the human eye. Finally, not wanting to risk prolonged exposure to invisible (but eye-damaging) laser beams, we did not perform range experiments with this setup. However, given the vulnerability of many microphones to IR as shown in Section 4.3.5.3, we conjecture that results similar to those obtained in Section 4.3.2 could be obtained here as well.

4.5 Future Directions

After a thorough investigation of LSI on MEMS microphones and the effects on VCSs, there is a significant amount of future work required to develop defenses and find vulnerabilities within other devices.

4.5.1 Recommendations for Defenses

The experiments and models developed for laser signal injection in MEMS microphones have shown that VCS devices are vulnerable to voice command injection. Future devices should be designed in a way to reduce or remove these capabilities.

Many defenses have been suggested to prevent these attacks. These can be detection mechanisms or mitigation techniques and generally fall into two categories: system-level defenses and sensor-level defenses.

4.5.1.1 System-Level Defenses

Some defenses would require changes to the software of vulnerable devices, potentially allowing for an update to defend against attacks. Some of the current generation of VCSs attempt to protect unauthorized execution of sensitive commands by requiring additional user authentication steps. For phone and tablet devices, the Siri and Alexa apps require the user to unlock the phone before executing certain commands (e.g., unlock the front door, disable the home alarm system). However, for devices that do not have other forms of input besides the user's voice (e.g., voice-enabled smart speakers, cameras, and thermostats) a digit-based PIN code is used to authenticate the user before critical commands are performed. This additional layer of authentication can be effective at somewhat mitigating the attack.

Unlike some other injection attacks, the *Light Commands* threat model lacks proper acoustic feedback, which makes it difficult for the attacker to pass any sort of liveness checks or continuous authentication methods. These can be as primitive as asking a user simple questions before performing a command, or as sophisticated as using data from different microphones [116, 117, 118], sound reflections [119], or other sensors [120] to verify that the incoming commands were indeed spoken by a live human.

As another potential software update, manufacturers can attempt to use sensor fusion techniques [121] in the hopes of detecting light-based command injection. More specifically, voice assistants often have multiple microphones, which should receive similar signals due to the omnidirectional nature of sound propagation. Meanwhile, when the attacker uses a single laser, only one microphone receives a signal while

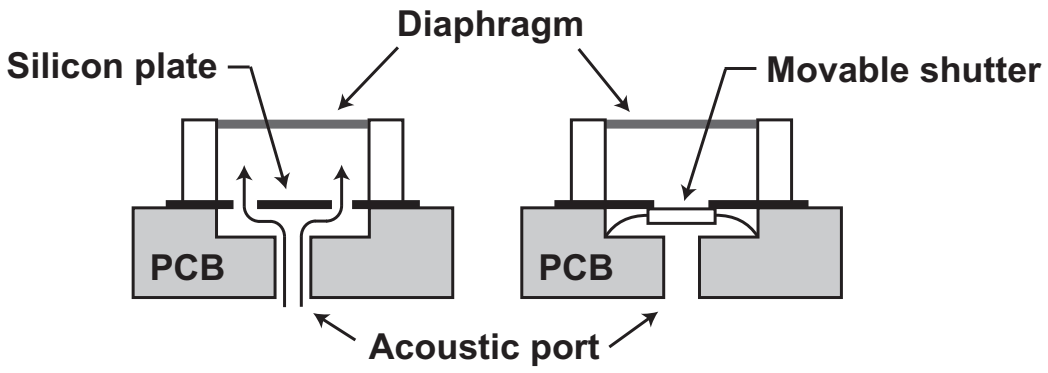

Figure 4.13: Designs of MEMS microphone with light-blocking barriers [1]

the others receive nothing. Thus, manufacturers can attempt to mitigate the attack presented in this paper by comparing signals from multiple microphones, ignoring injected commands using a single laser. However, attackers can attempt to defeat such comparison countermeasures by simultaneously injecting light to all the device's microphones using multiple lasers or wide beams.

4.5.1.2 Sensor-Level Defenses

The best way to protect all future MEMS microphones is to reduce the amount of optical energy that can enter the package of the microphone. This can be done by inserting barriers that will diffract, reflect, or block the straight optical path but allow sound to travel around it. This can be done at a system level with waveguides or light-blocking meshes, but this also can be accomplished at the device level with special light-blocking structures.

For example, the designs in Figure 4.13 have a silicon plate or movable shutter, both of which eliminate the line of sight to the diaphragm [1]. It is important that these barriers be constructed with materials that can block the light of a wide range of wavelengths, especially IR. While effective towards all laser effects, this blocking strategy often results in an inherent trade-off with acoustic sensitivity, as optical barriers are often acoustic barriers as well.

Beyond blocking light from entering the microphone entirely, the next best recommendation is to reduce the photovoltaic signal. This can be done by improving the coverage and optical properties of the glob top already used in MEMS designs for environmental protection. Our investigations show that there are gaps in the protection of the glob top where a laser signal can influence sensitive junctions on the ASIC, especially with IR light. Beyond this, devices such as the Vesper piezoelectric microphones are inherently more resistant to PV effects because the MEMS structure effectively blocks nearly all incoming light. This is not the case for thin silicon structures that will be partially transparent to a wide range of super-bandgap light wavelengths.

Besides preventing the photovoltaic effect, it is important to build MEMS designs that are resistant to thermoelastic bending. In our experiments, we found that several of the microphones only exhibited a very small thermoelastic bending signal, while some devices such as the SPU0410 had a significantly stronger bending signal. To prevent thermoelastic and plasmaelastic bending, it is important to develop MEMS designs and processes that reduce inherent thermal and electronic stress gradients within the diaphragm. As much as possible, the diaphragms should have symmetric thermal, mechanical, and electronic properties through its thickness.

The thermal diffusion effect is probably the most challenging one to attenuate, as by its nature, microphones require the diaphragm to be in contact with the air. Tuning acoustic parameters such as increasing the volume of the back package will attenuate the signal, but it will also affect the response of the microphone. The only other ways to reduce the TD signal is to reduce the temperature of the diaphragm by reflecting optical energy away or improving thermal connections to transfer heat away from the diaphragm.

Finally, this research provides several potential mechanisms to detect laser signal injection. From a system level, it may be possible to use signal processing of the

microphone signal to detect the low-frequency bias as an indication for a laser signal injection attack. Several microphones mounted on the same system can be used to check the validity of the incoming signal, potentially even using the unique phase responses to detect the presence of an attacking signal. On a device level, simple temperature sensors or light sensors can be intentionally designed into the MEMS or ASIC structure to indicate the presence of a strong light source. If the attacking signal can be detected, it can greatly improve the security of the systems using these microphones.

4.5.1.3 Hints towards Vulnerabilities in other Sensors

While this work breaks down the physical phenomena that lead to vulnerabilities within MEMS microphones, the phenomena that we investigated are not limited to MEMS microphones. This research provides indications of vulnerabilities in other sensors.

Any MEMS device that has an opening to allow light to enter the package is potentially vulnerable to photoelectric signal injection via laser. As MEMS structures are often designed with silicon, concentrated IR light can potentially transmit through any MEMS structure and affect ASIC circuitry. This could be a concern for any device designed to interact with an external fluid, such as MEMS ultrasonic sensors, pressure sensors, humidity sensors, or chemical sensors.

Beyond this, any sensor that uses the motion of a mechanical structure that is exposed to the environment is potentially vulnerable to photoacoustic signal injection. This includes conventional microphones, ultrasonic sensors, and pressure sensors. Our work discusses the many potential ways that this photoacoustic signal can be generated, all of which will be highly dependent on the structure and materials of the mechanical structures used in these sensors.

Finally, any sensor that uses the motion of a mechanical structure within an en-

closure of air or another gas may potentially be vulnerable to photoacoustic injection via thermal diffusion effects. Thermal Diffusion only requires that air within the enclosure be heated periodically, which can potentially be accomplished by heating the sensor package itself instead of any exposed MEMS structure. This would include MEMS accelerometers, gyroscopes, magnetometers, and oscillators that have movable MEMS structures within an enclosure of gas.

CHAPTER V

Characterizing Laser Signal Injection on Space Systems

Space is an emerging commercial critical infrastructure that requires extensive security analysis and protection [122]. As of 2023, over 2,000 small satellites have been launched already, and more are well on the way with the running total increasing almost exponentially [123]. Meanwhile, the number of observed satellite attacks also increased proportionally [124]. A large portion of these past incidents operated in conventional computer and information security domains such as software access controls and wireless communication protocols [124, 125]. Similarly, academic research in space security had mostly focused on the wireless communication links of satellites [126]. Protection of these digital system components alone, however, is insufficient because as cyber-physical systems, satellites feature analog interfaces such as sensors whose output can have direct influence or control over the space system's behaviors.

Previous military and aerospace research has already verified that physical signals such as lasers can jam or damage sensory components of space systems, compromising the availability of satellite sensors. Meanwhile, the research presented in this chapter[1] shows that specially modulated laser signals can induce controlled outputs from

[1]B. Cyr, Y. Long, T. Sugawara, and K. Fu, "Position Paper: Space System Threat Models Must

sensors, compromising the integrity of some sensor-reliant systems on earth. If these two lines of research are integrated, it leads to the natural and intriguing follow-up question: to what degree can laser signal injection also be used to compromise the integrity of satellite sensor readings? Sensor integrity attacks can usually be more stealthy and provide more malicious control over the target systems. There remains a gap of knowledge for published studies analyzing or defending satellite systems for sensor integrity vulnerabilities.

The purpose of this chapter is to give a preliminary investigation of laser signal injection attacks on space systems. It will answer the following questions:

- What potential ways can space systems be exploited by LSI?
- What capabilities exist to perform an LSI attack on sensors in space?
- What are the potential consequences of LSI on space systems?
- What open problems and research directions exist for LSI on space systems?

5.1 Sensors in Space Systems

Satellites almost always employ a variety of sensors to measure their environment and determine the current state. Space systems consist of many different subsystems that rely on sensor data to fulfill mission requirements. The common subsystems with sensor components are the Attitude Determination and Control Subsystem (ADACS), Electrical Power Subsystem (EPS), Communications Subsystem, Thermal Control Subsystem, Propulsion Subsystem, and the Payload.

Within the ADACS, many different sensors are used to determine its attitude (i.e. its position and orientation) to provide feedback for control maneuvers. For a space system to determine its attitude, at least two vectors need to be related from

Account for Satellite Sensor Spoofing," in *SpaceSec23*, Feb. 2023.

the "body-fixed frame" of the satellite's perspective to the "inertial frame" of its relation to other objects in orbit. For low-cost satellites, the two vectors often come from a measurement of the earth's magnetic field via a magnetometer, and from a measurement of the direction of the sun with one or more sun sensors [127]. These sun sensors are often one or more photodiodes that use the photovoltaic effect to find the sun. In more advanced satellites, the attitude vectors can be determined using cameras that either perform horizon sensors to locate the Earth's horizon [128] or star tracking to locate star constellations [129]. Beyond these devices, various other sensors are often used to support attitude determination, including IMU sensors to measure acceleration and angular velocity.

For the EPS, many sensors are used to measure battery status, power usage, and deployment state. Beyond these electrical sensors, photovoltaic (PV) cells, which nearly every satellite relies upon to obtain energy, form a mechanism to enable signal injection into the EPS. Satellites often rely upon different algorithms to perform Maximum Power Point Tracking (MPPT) and obtain as much energy from incoming light as possible. MPPT algorithms rely upon the signal from the solar cells to perform their operations. Beyond this, solar cells can even be used as a coarse sun sensor during some operations [130].

Most conventional designs of the Communications Subsystem rely upon some form of radio communication, but there has been increasing interest in the use of light as a medium for communication. These optical communication systems are used instead of conventional radio systems due to their relatively low energy divergence and high directivity, allowing for high bandwidth communication at lower mass and power budgets [131]. The main limitation of this communication is a requirement for a line-of-sight (LOS) between transmitter and receiver to perform communication, which can make it more difficult to construct a link. Because of the advantages of optical communications, significant effort has been put into developing the technology

of ground-to-space and space-to-space laser transmitters and receivers [132]. These receivers can themselves be considered sensors, relying on Avalanche Photodiode (APD), CMOS cameras, or other photo-sensitive technologies [132].

Another sensor-reliant subsystem is the Thermal Control Subsystem. Due to the extreme conditions of space, keeping satellite components within the necessary parameters can become a challenge. Temperature sensors are often utilized throughout the satellite to monitor the thermal conditions. A satellite will often make control decisions to radiate, retain, or generate heat depending on these sensors.

The Propulsion Subsystem is present in many space systems as a way to perform precise attitude adjustments. While technically part of ADACS, the complexity of these devices often allows them to be considered a separate subsystem. Propulsion devices often rely on sensors such as pressure sensors to determine how much propellant remains ready for use.

Finally, the primary mission objective of a space system is to provide capabilities for a payload consisting of sensors designed to carry out a specific mission. The primary payload of many satellites is often optical sensors. This is often in the form of visible-light and infrared cameras to monitor conditions of the earth, hyperspectral cameras [133] to gather scientific data or photodiodes for sensing nuclear detonations [134].

While all of these sensors have been used to develop reliable ways to monitor the state of a spacecraft and perform its mission, very little has been done to understand their vulnerabilities to cyber-physical threats.

5.2 Related Work

There is already a growing concern over potential cybersecurity attacks on the software, networks, and communications systems used by these satellites, and even a confirmed attack on satellite operators [135]. Some researchers have been warning

of the dangers of ignoring security on modern space systems [136, 137]. Because of this, recent works at small satellite conferences have been publishing works on ways to improve network security and trust within the small satellite community [138, 139]. Some companies have also started offering specialized consultation in aerospace security [140]. Most of this research is new and predictive rather than empirical, due to the difficulties in experimenting with space systems. Several surveys on the topic have been written [124, 125, 126, 141] to give a more comprehensive overview.

Beyond these software and networking attacks, a substantial body of literature from the aerospace and optical engineering communities has verified that physical signals in the form of lasers and electromagnetic waves can be used to compromise the availability of satellite sensors by damaging or jamming sensor-related functions. For example, one work describes how commercially available pulsed ground-based lasers could be used to damage the solar arrays on some satellites [142]. A similar work found that airborne lasers and potentially ground-based lasers can damage photodetectors in a generic space telescope in geostationary earth orbit [143]. Another research work showed both theoretically and experimentally that an 8 kW laser can jam a MSTI-3 satellite's photodetectors from 5 km away [144]. Further analysis confirmed that high-energy laser can be used for jamming or blinding space-borne photoelectric sensors, destroying satellite solar cells, and destroying satellite thermal control systems [145]. In view of these existing attacks, some military efforts have also been spent on developing techniques for detecting and warning laser-based attacks [146]. However, research in this area did not consider more advanced sensor integrity attacks that aim to control sensor readings more stealthily with LSI.

5.3 LSI Attacks on Space Sensors

Laser signal injection attacks on space sensors are attacks on sensor data integrity. An example of this attack is illustrated in Figure 5.1. Rather than simply stopping a

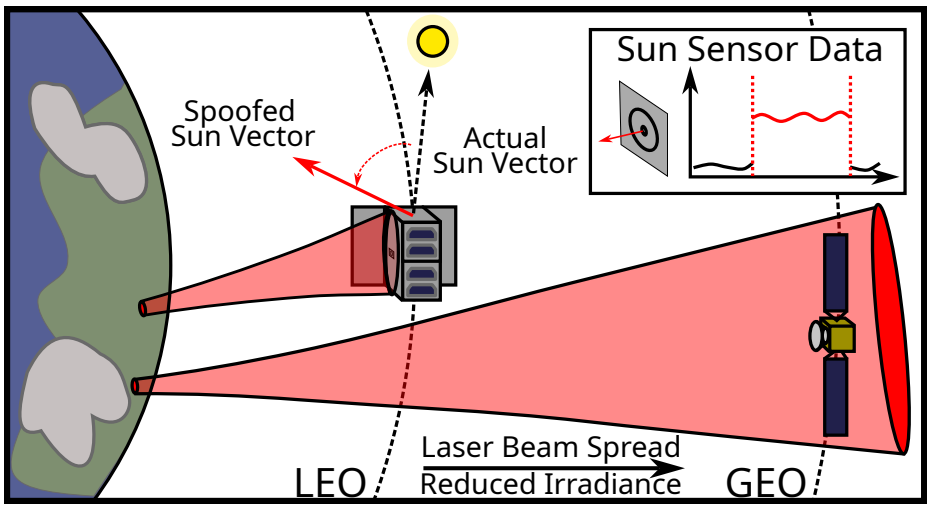

Figure 5.1: An attack scenario where a high-powered laser is used to inject a signal into sun sensor data, spoofing a new sun vector. This is more effective at Low-Earth Orbit, as increased distances reduce the laser signal irradiance.

sensor from receiving data, a signal is injected into the system, leading to attacker-controlled data that can be leveraged to control the system. This type of attack has significant challenges within the context of space systems, so very little is understood about the true potential for LSI vulnerabilities. The purpose of this section is to perform a preliminary investigation into LSI attacks on space systems to determine the potential for vulnerabilities within these devices.

In all of these investigations of LSI in space systems, the three primary threat models to consider are:

1. **Ground-to-Space** [142]: An adversarial laser signal is fired from a ground-based system through the atmosphere. There are significantly fewer restrictions on power and optics.

2. **Air-to-Space** [143]: An adversarial laser signal is fired from a high-altitude aircraft to bypass most of the atmosphere. In this model, there are more constraints to what powers and optical conditions are possible.

3. **Space-to-Space** [147]: An adversarial laser signal is fired from a space system

Table 5.1: A summary of potential LSI attack scenarios against space sensors

Sensor Type	Associated Satellite Sub-systems	Example Attack Scenario	Selected References
Star/horizon Tracker	ADACS	Spoofing a star formation or horizon to change perceived orientation	[4, 80, 78]
Light Sensors	ADACS	Spoofing or changing a sun vector to cause incorrect ADACS decisions	[148, 83]
Inertial Measurement Unit	ADACS	Light-generated signal to spoof angular inertial changes	[149, 64, 63]
Photovoltaic Cell	EPS	Signal Injection into the power system to create faults or reduce efficiency	[148, 150]
Temperature Sensors	Thermal Control	Localized heating of an area, resulting in heating or attitude shifts	[151, 62]
Pressure Sensors	Propulsion	Laser-generated signal to spoof changes in propellant density	[85, 64]
Camera	Payload	Inject controlled patterns into images that hide or alter real objects	[78, 80, 152]

to influence other systems in space. This threat model contains the highest amount of constraints on power and optics.

Each of these threat models has different considerations, but all have the potential to perform LSI on space systems.

5.3.1 An Overview of Potential Attacks

It is important to investigate commonly used sensors and define the potential attack surface for sensor spoofing attacks. In investigating the attack surface, we are not considering conventional attacks on communications, but we are instead focusing on LSI attacks on sensors used in all other subsystems. A summary of the sensors considered in the attack surface is listed in Table 5.1.

5.3.1.1 Attitude Determination and Control

The Attitude Determination and Control Subsystem (ADACS) in a space system is responsible for measuring and adjusting the attitude (orientation) of the entire system. The ADACS is critical for many orbiting devices, as precise pointing of sensing instruments and antennas is required to fulfill mission requirements. The subsystem relies on an automated control loop of several sensors to control attitude, which makes it an attractive target for sensor spoofing.

Star/Horizon Trackers. Star trackers and horizon trackers are both camera systems designed to determine the satellite's attitude by locating fixed references to determine the relative orientation of the system. For star trackers, an algorithm

matches the stars to a known database of constellations. For horizon trackers, an algorithm locates the horizon of the Earth as the fixed reference. Since these sensors are simply cameras, an incoming laser signal will add additional information to the image that is parsed by the underlying algorithms. By exploiting features of the camera such as frame rate, a rolling shutter [80], or lens flare [78], an attacker may exhibit a level of control on the output of the trackers without the faults generated by a simple jamming attack. Depending on the attacking signal and the algorithms, the trackers can report incorrect orientations to the ADACS controller, and cause a change in satellite attitude. This will reduce system performance or prevent the system from accomplishing its mission.

Light Sensors. Light sensors such as sun sensors and bolometers are photosensitive components that are mounted in a way to give an estimation of the location of the sun or earth relative to the body frame of the system. They often consist of a set of photodiodes, 2D photodiodes, or photoresistors mounted in a way that visible or infrared light from the sun or earth will hit different photosensitive components at different orientations [153]. By comparing the signal between the light sensors, a rough vector to the sun or earth can be computed and used for attitude determination. Spoofing attacks on light sensors have already been demonstrated [148, 83], which suggests a vulnerability to spoofing is likely. By spoofing the light signal, an attacker can change the measured light vector and gain some control over the attitude control.

Inertial Measurement Units (IMUs). IMUs are a collection of sensors meant to determine the inertial changes to the body of the system. In the case of orbital systems, a gyroscope and magnetometer are often employed in tandem to measure angular inertia. Conventionally these sensors were built mechanically or optically with large parts, but more recently smaller satellites have been relying more on MEMS components. Due to their smaller size, MEMS sensors inherently have less inertia and

more susceptibility to injected signals. Research on laser-based attacks on MEMS sensors is limited [149, 85], but the potential exists that changes to the thermal or mechanical state of the system can induce changes to the output of these devices. If an attacker can affect the output of these sensors, it would give them significant control over the attitude of the system.

5.3.1.2 Electrical Power Subsystem: Photovoltaic Cells

The Electrical Power Subsystem (EPS) of the space system is responsible for providing the necessary electrical energy to the rest of the components. Nearly all systems in orbit rely on energy generated from photovoltaic (PV) cells that collect light energy from the sun. These photovoltaic cells are often used in conjunction with special circuitry to perform maximum power point tracking (MPPT) control algorithms to maximize the energy output from the PV cells [154]. Since PV cells are designed to capture as much light as possible, they are particularly vulnerable to laser signal injection attacks. PV cells are sometimes used as coarse sun sensors for attitude determination [130], leading to the same sensor spoofing vulnerabilities as light sensors [148]. An attacker can also use the PV cells to inject a signal into the power system directly. Depending on the design of the EPS, a number of power injection attacks may be possible, similar to the ones used in [150]. Beyond this, the PV cells and subsequent power distribution components produce a significant amount of electromagnetic noise [155], which can potentially be leveraged to disrupt measurements or inject signals into other parts of the system.

5.3.1.3 Thermal Control: Temperature Sensors

The thermal control subsystem is critical in space, where extremes in temperature can push components out of the operating ranges and risk component failure. Various temperature sensors are used to measure the temperature distribution in the space

system, allowing thermal control to use heaters or request attitude adjustments to ensure safe temperature ranges. Temperature-critical systems have been shown to be vulnerable to sensor spoofing [151, 62], and we expect space systems to be similar. As heating is a primary mechanism by which light will interact with the space system, the temperature sensors will be inherently vulnerable. Spoofing attacks could lead to excess power usage, attitude shifts, or system faults caused by overheating, as it is difficult to cool the system efficiently.

5.3.1.4 Propulsion: Pressure Sensors

Many space systems require propulsion subsystems to adjust orbits or attitudes. These systems function by storing gas propellant that can be fired in short bursts when needed. Pressure sensors are used to measure the status of the propellant and report to the rest of the system. If a laser signal can heat the propellant, generate a photoacoustic signal [64], or exploit photoelectric effects [85], it could potentially spoof incorrect propellant status to cause control errors or misfires.

5.3.1.5 Payload: Optical Sensors

The primary payload of many satellites is often optical sensors. This is often in the form of visible-light cameras, infrared cameras, hyperspectral cameras [133], or photodiodes for sensing nuclear detonations [134]. These sensors would be particularly susceptible to an adversarial laser signal, as any incoming light will be focused by a lens upon the optical sensor. At low irradiance levels, this will simply be a noise source localized to the set of pixels describing the location of the source of the attacking signal. At higher irradiances, light reflections and scattering within the optics will lead to lens flare, creating noise on a much larger part of the image [156]. While data from these sensors are not usually critical for the system to function, future applications using automated computer vision systems could be vulnerable. This is

seen by example within the autonomous vehicle community, where computer vision systems are susceptible to sensor spoofing with lasers through various mechanisms [80, 78].

5.3.2 Characterizing Attacker Capabilities

Laser signal injection attacks on space sensors require special attacker capabilities. While many of the potential threat vectors described in Section 5.3.1 may be possible in controlled laboratory environments, the extreme distances and conditions of space systems require significant attack capabilities. The purpose of this section to describe some of the capabilities that will be necessary to perform LSI on space sensors.

5.3.2.1 High-Power Laser Capabilities

While space technology has become increasingly dense and closer to Earth, optical technology has been improving to provide higher power over longer ranges. State-sponsored laser research into directed energy weapons (DEW) has led to many new technologies for long-range, high-powered lasers [157]. In the United States, programs such as ALPHA and MIRACL [158] used megawatt class hydrogen-fluoride lasers with beam directors a few meters in diameter to investigate anti-satellite (ASAT) capabilities. Both Russia [159] and China [160] are developing laser anti-satellite technologies.

There has also been growing research and development into fiber laser systems, which use doped fiber optic cables as a gain medium. These devices are stable, have higher beam qualities, and can produce several kilowatts of power [161]. This has led to the development of fiber-laser technology with beam combination optics for use in DEWs, such as the 33kW Raytheon Laser Weapon System (LaWS) [157], the 50kW DEM-SHORAD [162], and the 100kW Dynetics-Lockheed HEL TVD [163]. Fiber lasers have also enabled companies to build kilowatt-class fiber laser systems

for welding and cutting, increasing the availability of high-power lasers [164]. We expect to see the continued development of laser technology that will make lasers high-powered and easier to obtain, increasing the capabilities of an attacker to intelligently inject signals.

5.3.2.2 Effective Range

While earthbound sensor spoofing attacks have only been demonstrated to work at ranges less than 100 meters, we have reason to believe that high-powered lasers can be designed to spoof at much farther ranges. The primary parameter that will enable attacks on the sensor integrity of space systems will be the irradiance (power density) of the attacker's laser signal at the vulnerable component. Space is large, with distances in the tens of thousands of kilometers just within the space systems in Earth's orbit. Because of this, it is important to understand how electromagnetic energy diffuses at long distances.

The fundamental limiting factor for the effective range of any laser beam is diffraction. This serves as a hard limit for possible laser attacks. As described in Section 2.2.3, the irradiance I over distance z for a collimated, diffraction-limited Gaussian laser beam will have the following relationship [10]:

$$I(r,z) = \frac{2P_0/(\pi w_0^2)}{1+(z/z_R)^2} \exp\left(\frac{-2r^2/w_0^2}{1+(z/z_R)^2}\right) \tag{5.1}$$

$$z_R = \frac{\pi w_0^2}{\lambda} \tag{5.2}$$

where P_0 is the total optical power of the beam and z_R is the Rayleigh length defined from the wavelength λ, the beam quality M^2, and the beam waist w_0. In the far-field case ($z \gg z_R$), the irradiance of the laser follows an inverse square law, but larger beam sizes and shorter wavelengths at the transmitter will greatly increase the effective range.

5.3.2.3 Timing and Modulation

The primary difference between a spoofing attack and a denial-of-service attack on a sensor is the timing and modulation of the injected signal to achieve a stealthy and effective attack. For previous works investigating sensor spoofing attacks, special care had to be taken to inject spoofed signals rather than simply overwhelming the sensor with noise. Sensor spoofing attacks on space systems will be no different. Developing appropriate modulation techniques with lasers will be a challenge, as high-power lasers have technological limitations on the precise control of the output irradiance. For example, pulsed lasers are often used to deliver extremely high power in short pulses, but are often limited in repetition rates, as it takes time to cool and charge the gain medium.

5.3.2.4 Angles and Aiming

One of the hardest challenges to overcome in performing sensor spoofing attacks is aiming the beam at appropriate angles. Lasers can only attack sensors within line-of-sight, preventing attacks on sensors protected by the earth or the body of the satellite. This is especially important for earth-to-space attacks on low earth orbit systems, where transits across the sky last on the order of minutes. Beyond this, there is a fundamental trade-off between smaller beams with higher irradiances and the precision required for aiming. The challenges of tracking and aiming the beam for a consistent spoofing attack will be considerable.

5.3.2.5 Atmospheric Disturbances

A limiting factor for injection attacks from ground-based and air-based attack scenarios lasers is disturbances caused by firing a laser through the atmosphere. This area has been greatly studied to improve capabilities in astronomy and satellite communications, but it still is a significant challenge to any long-range, laser-based attack.

In the atmosphere, four mechanisms will affect lasers: scattering, absorption, turbulence, and thermal blooming [165, 166]. Each of these mechanisms will reduce and add randomness to the irradiance at the target system.

Scattering and absorption result in a reduced beam irradiance as light energy is lost in accordance with the Bouguer-Beer-Lambert (Section 2.2.4. This puts restrictions on potential wavelengths used in ground-based attacks, as the atmosphere selectively absorbs various wavelengths in differing amounts. Turbulence is the change in the beam induced by the refraction of light through turbulent air currents. This results in an increased beam size, random movement in the beam (beam wander), and random variations in irradiance throughout the beam at the target (scintillation). Finally, thermal blooming caused by the heating of air on a stationary transmitter causes the beam to refract, further reducing irradiance for a very high-power laser.

The atmospheric disturbances to a propagating light have been modeled by many different groups to enable ground-to-space laser communication and to minimize optical distortions while taking astronomical measurements. One such model for the mean irradiance at the target ($I_s(z)$) of a ground-to-space laser with a large beam waist is [166]:

$$\langle I_s(z) \rangle = \frac{I(z)}{1 + (2\sqrt{2}w_0/r_0)^{5/3}} \tag{5.3}$$

where $I(z)$ is the irradiance of the Gaussian beam without the disturbance, w_0 is the beam waist, and r_0 is the atmospheric coherence width or Fried's parameter. The atmospheric coherence width is a special parameter that indicates the turbulence of the atmosphere at the firing location. It can be modeled by:

$$r_0 \approx \left[\frac{16.7 \int_{h_0}^{H} C_n^2(h) dh}{\lambda^2 cos(\zeta)} \right]^{-3/5} \tag{5.4}$$

where λ is the light wavelength, ζ is the zenith angle describing how long the laser is traveling through the atmosphere, and C_n^2 is the atmospheric turbulence strength at

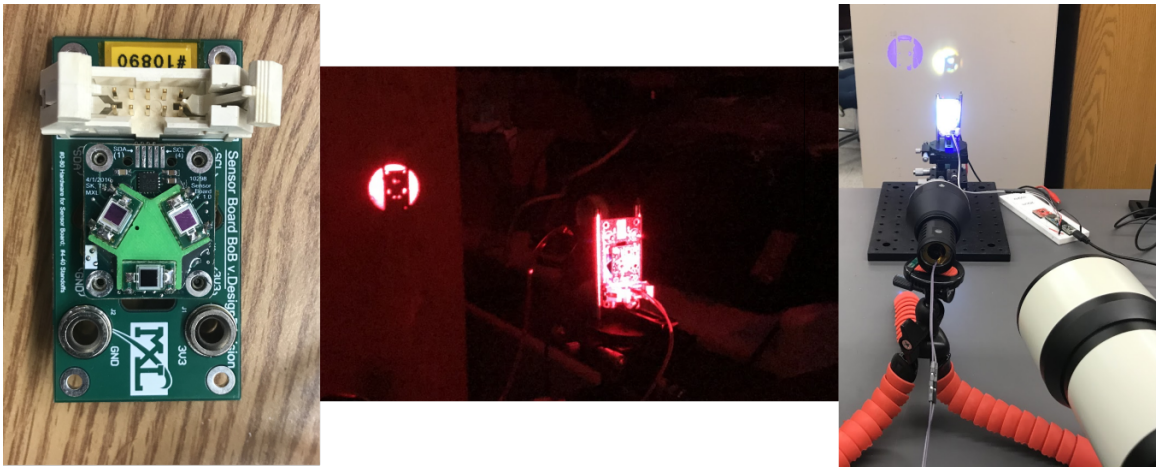

Figure 5.2: (Left) Triclops Sun Sensor used in UM Cubesat program. (Middle and Right) Using lasers and lights to simulate a laser-based injection attack

each altitude from the height of the laser (h_0) to the height of the target (H). While C_n^2 is complex and highly random, these models indicate the additional challenges in firing a laser through the atmosphere. Short wavelengths, which have less diffraction as they propagate, are actually scattered more readily by air, reducing the irradiance at the target. Atmospheric conditions throughout the air above the attacking laser will greatly affect the results and make it difficult to control the irradiance at the target.

The effects of atmospheric disturbances greatly reduce the attacker's control over the attacking signal, and this may require special techniques such as adaptive optics [167] to overcome this limitation.

5.3.3 Case Study: Sun Sensors

While the potential threat vectors and considerations are described across many different sensors in space systems, a preliminary investigation was performed to characterize the attacker capabilities of LSI on sun sensors. Working with the Michigan eXploration Laboratory (MXL) [168], a "triclops" custom sun sensor was obtained for experiments. The triclops is a sensor with three sensitive photodiodes that are at

slight angles away from each other (see Fig. 5.2-Left). Using a set of transimpedance amplifiers, the triclops was developed specifically to locate the direction of the sun by measuring the differences in photo-generated current on the faces of the three photo-diodes. These currents are converted to three voltages and measured using an ADC, and the signal was used in a satellite attitude determination algorithm to determine the sun vector. The first set of experiments to explore the cyber-physical capabilities of an attacker was to see how much the sun vector could be affected by incoming laser light.

To do this, a simple experiment was set up to model a laser-based injection into the triclops sensor. The goal of the experiments was two-fold. First, the experiments measured the sensitivity of the triclops to various wavelengths and incident angles of incoming laser light. Second, the experiments gave some initial results into the optical power density that would be required to generate a meaningful change to the output signal. In the experiment, three different laser diodes were used: a 450 nm Osram PLT5 450B blue laser, a Thorlabs L520P50 520nm green laser, and a 638nm Thorlabs L638P150 red laser. These diodes were mounted onto a telephoto lens to expand the collimated laser beam, which was necessary to create a uniform optical irradiance (i.e. power density) over the surface of the device. A ThorLabs LDC205C laser driver was used to control the diode current and output power, and a ThorLabs PM100 power meter was used to measure irradiance. As a model for the sun, an Acebeam W30 flashlight was used to provide high-irradiance, broad-spectrum light. A ThorLabs LDC240C was used to control the irradiance of the flashlight. The flashlight was placed perpendicular to the triclops surface, while the laser was placed at a 10° angle (Fig.5.2-right).

Figure 5.3 shows the results of these preliminary experiments. They show that the photodiodes used in this device are significantly more sensitive to green light than other colors of light. Beyond this, it takes significant irradiance to produce large

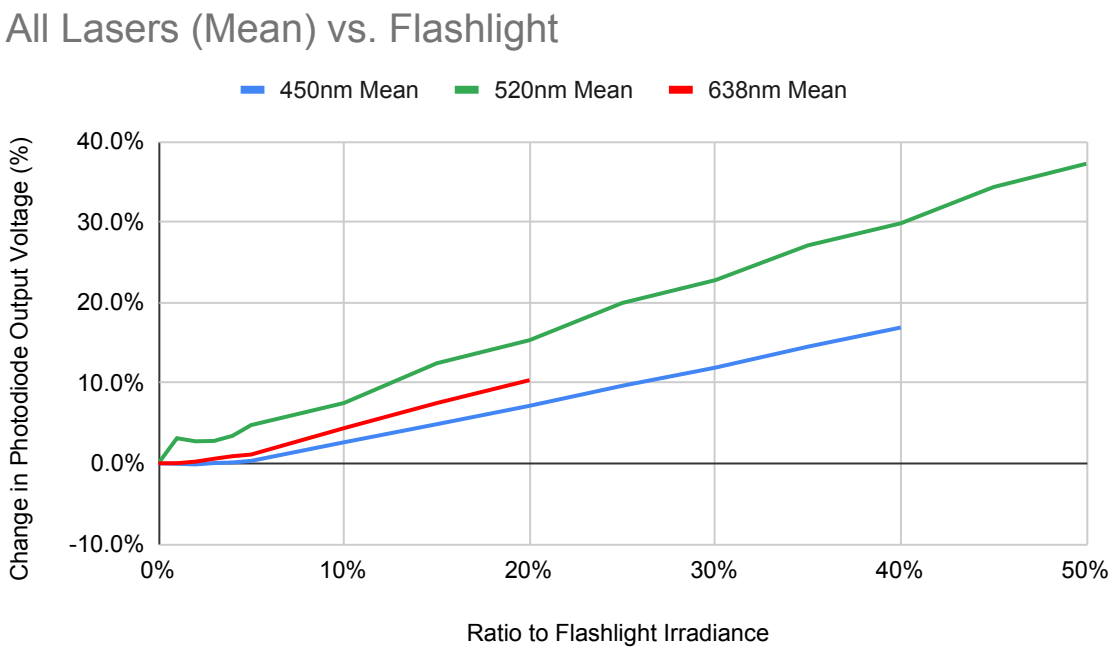

Figure 5.3: A comparison of the effects of different colors of laser light on a sun-exposed triclops.

changes in the output signal. In order to perform light signal injection into these photodiodes when the sun is directly on them, there would need to be irradiances comparable to the sun, which can be difficult at distances in space.

This one experiment is just a preliminary investigation into a single sensor used in space systems. Further investigations need to be done just to discover how realistic a threat scenario can be developed.

5.4 Consequences on Satellites

While this work is a preliminary investigation into LSI on satellites, it is important to understand some potential consequences on current and even future space systems. The primary concern of any LSI attack on a satellite is the potential for a stealthy mechanism to generate a loss of performance on the targeted system. This can be temporary, where the LSI causes the system to become incapable of performing its

primary mission for a limited amount of time before returning to normal operation. Or the loss of performance can be permanent, where the system is damaged in a way that lowers its capabilities or even causes mission failure.

This loss in performance can occur in several ways. First, attacker influence over the attitude can potentially lead to poor pointing accuracy, which can be vital for a system to carry out its mission. In some instances, a significant shift in pointing accuracy could lead to the damage of sensitive components by pointing them toward the sun. In other cases, a shift in pointing accuracy may prevent the use of the communication subsystem. Second, as systems become more automated and designed to capture specific phenomena, LSI on various optical sensors may cause them to have poor accuracy in capturing the phenomena. This would be especially concerning for high-priority sensors such as the ones used to detect nuclear detonations. False positives and false negatives could lead to many problems. Finally, an LSI may harm performance by causing an excess amount of energy to be expended. Satellites run on limited power budgets, and any LSI-induced action will require the expenditure of the limited energy, potentially preventing other components from using that energy.

Ultimately the key feature of LSI is its stealthiness. In the context of space systems, a jamming attack causing a denial-of-service may be easily detected, immediately suggesting the presence of an attacker. But an LSI attack can be used to shift data measurements, leading to errors that can be attributed to a faulty sensor. This has important ramifications within space systems, where a significant amount of money and concern is placed on the security of these devices. A measured laser attack on a system by another entity generates significant concern. A successful LSI attack, however, simply seems like a system failure due to a faulty sensor, causing the attack to remain undetected.

5.5 Future Directions

The information in this chapter is a preliminary look into this research area, with the goal of setting foundations to lead to future works. Here are some recommendations for future defenses and areas of research to be considered.

5.5.1 Recommendations for Future Defenses

Due to the challenges of LSI in space systems, the primary way to prevent LSI attacks is to have reliable ways to detect the presence of an attack. This strips a sensor integrity attack of its primary advantage of being stealthy and allows a space system to perform operations to go into a safe mode. There are two primary ways to detect the presence of LSI attacks: anomaly detection and the use of optical sensing capabilities.

Updates to software to register sensor anomalies would be an important step in detecting LSI attacks. While a general defense to any injection attacks on sensors, anomaly detection in space systems will be very effective, as the state of a controlled system will change slowly and can be modeled reliably. This means that defenses that detect sudden changes in state such as the one proposed by Choi et. al. [91] have the potential to be very effective.

Beyond anomaly detection via software, there are some simple changes to hardware or the EPS that could be used to detect LSI by the incoming light. Due to the ranges and disturbances involved in an LSI attack on space, it is very difficult to focus a beam on any particular component of the target. Diffraction will spread the beam over the whole target, and some simple components to measure this signal can be used to detect the presence of an attack. This can be accomplished with sun sensors or solar cells already on the device, as a sudden increase in the optical irradiance striking the system will be an indication of an attack.

5.5.2 Open Problems

Previous research has shown that absolute trust in sensor data creates susceptibility to sensor spoofing attacks. Space systems are expected to exhibit similar vulnerabilities as technology develops and space becomes more accessible. The purpose of this work is to encourage of research environment to investigate new threat models that exploit satellite sensor LSI, so that future space systems can be designed to protect against these threats. To accomplish this, several open problems will need to be investigated:

1. Models and simulations will be needed to determine attacker capabilities and limitations in satellite LSI.

2. Controlled experiments will be necessary to measure the vulnerability of sensors used in space systems to LSI.

3. Test beds need to be developed to better determine the full system consequences of satellite LSI.

4. Methods should be developed to provide forensic analysis in the case of an LSI attack against sensors in space systems.

5. Mechanisms should be investigated to reduce the risk of LSI in systems already deployed in space.

6. Robust mechanisms need to be developed to detect LSI in all classes of space systems.

Future research into these areas will help ensure LSI attacks do not become a concern in the future of increasing capabilities in space.

CHAPTER VI

Conclusion

Modern cyber-physical systems are vulnerable to laser signal injection attacks on the sensors they rely upon. These laser signal injection attacks exploit various physical phenomena, affecting sensors in ways that were never expected or intended. As cyber-physical systems become more prevalent, it is important for future systems to be aware that sensors often measure much more than they were designed to measure.

The purpose of this research is to perform a characterization of laser signal injection into three different cyber-physical contexts. In LiDAR systems used in autonomous vehicles, the characterization revealed an increase in attacker capabilities in comparison to previous works, leading to significant consequences on the safety of autonomous vehicles. In MEMS microphones used in voice-controllable systems, the characterization explored a brand new vulnerability to determine the key factors that lead to a vulnerability in devices such as smart homes. Within space systems, the characterization is an exploration into the security of sensor data integrity in space systems, highlighting the sensors and capabilities that should be explored in further investigations. These characterizations provide a way to better understand the potential vulnerabilities and consequences of these attacks, allowing future designers to assess risk and develop effective defenses.

Looking forward, the discoveries and research within this work provide many

opportunities for further advances in research. This future research would be necessary to ensure that future devices become resistant and remain resistant to incoming laser signals. There are three main branches to explore: the development of effective defenses, the discovery and characterization of other cyber-physical threats, and educational methods to increase awareness of signal injection attacks.

There is significant work to be done to ensure the devices within this dissertation remain secure against laser signal injection attacks. In the case of LiDAR within autonomous vehicles, the next steps would be to incorporate more context into object detection algorithms, make optical adjustments to the LiDAR to reduce the vulnerable field-of-view and investigate ways to randomize or encode the LiDAR firing sequence in a way that makes it difficult to an attacker to predict. In MEMS microphones in voice-controllable systems, the areas to investigate would be mechanisms to authenticate users by sound, the use of coverings or waveguides to block light, or developing microphones with protection to photoelectric and photoacoustic effects. For the sensors in space systems, efforts should be made to develop mechanisms to detect incoming laser signals and improve control systems to remain robust in the case of anomalous sensor data. Development and testing would be needed to determine which defenses will remain effective without compromising the performance of each of these systems.

Beyond defending these specific devices from laser signal injection, there is also significant work remaining to determine other vulnerabilities to laser signal injection. While this work was limited to three contexts, many other sensors and devices can be expected to have similar vulnerabilities to laser signal injection. More work needs to be done to explore other sensors and devices to determine new attacker capabilities to influence underlying systems. This includes optical sensors in various contexts, but also sensors such as MEMS sensors that are not inherently made to measure optical energy. There are many physical phenomena that light can use to interact with these

sensors that were not explored in this work but can potentially lead to vulnerabilities in other systems. As laser technology develops with increasing capabilities, there will need to be continued research to characterize attacks on these other sensors and devices.

Finally, in order to protect against laser signal injection attacks, there needs to be the development of methods, processes, and tools to educate and assist system designers in the defense of these attacks. This will include various ways to develop courses and other ways to teach about laser signal injection and sensor security in general. There are also research opportunities to incorporate sensor spoofing threat models into processes and metrics that are being developed to assist system designers to build secure systems. This research would also include the potential for the addition of these threat models into the tools that system designers use to break down the knowledge gap between how a sensor is expected to function versus how it physically measures the world. The education of these threat models and the encouragement of future designers to build secure systems is still a difficult problem to be investigated.

From the realm of LiDAR and autonomous vehicles, microphones in voice-controllable systems, and potentially other sensors used in the extreme conditions of space, light can be used to change a system's perception and influence control decisions. It is important to fully characterize the attacker capabilities and consequences of laser signal injection, as this knowledge can be used to guide future design decisions in the realm of cyber-physical systems. This thesis has contributed to closing the gap between expectations and the reality of sensor vulnerabilities, with the goal of inspiring further effort into designing safe and secure cyber-physical systems.

APPENDIX

APPENDIX A

Other Mechanisms for LSI in MEMS Microphones

In Chapter IV, there is an investigation into what is considered to be the most likely effects generating the output seen in laser signal injection into MEMS microphones. But there are a few other phenomena that can potentially influence MEMS microphones. For most devices, the contribution of these extra phenomena will be very small, but can potentially become a problem for highly-sensitive sensors or future microphone designs. The purpose of this appendix is to highlight these different effects and show why they are not considered a primary contributor.

A.1 Models for Other Effects on MEMS Microphones

To investigate these other effects, a simple model is developed for each phenomenon in order to explain why it has only a small contribution to the output signal.

A.1.1 Plasmaelastic Bending in Asymmetric Diaphagms

Within semiconductor materials, the photogeneration of electron-hole pairs will cause elastic deformations within the structure. This is due to the electrons jumping

from the valence bands to the conduction bands within the semiconductor material, changing the overall charge distribution within the crystal structure. For silicon samples, the crystal structure actually contracts in response to this change in charge distribution, opposing the effects of thermal expansion. This effect was first discovered by Gauster and Habing [22], and it came to be known as the concentration-deformation mechanism [23], electronic deformation [169], or the plasmaelastic effect [26]. For many MEMS microphones, the diaphragm is made out of doped polysilicon, which can have these plasmaelastic properties.

Plasmaelastic bending is caused by a moment generated as a semiconductor changes its volume in response to the generation of charge carriers. This effect was also described in Todorović et al. [25] in a similar manner to Thermoelastic Bending, only with the relevant parameters related to the minority charge-carrier density instead of temperature. Similar to temperature, any generated charge carriers will diffuse almost immediately throughout the thickness of the MEMS structures. Because of this, the charge carrier generation can be approximated as uniform throughout the semiconductor portion of the MEMS structures. The excess minority charge-carrier density (Δn) in the diaphragm can be described as:

$$\Delta n \approx \frac{\lambda I_A}{hcL_s(j\omega + 1/\tau)} \tag{A.1}$$

where L_s is the thickness of the semiconductor portion of the diaphragm, and τ is the minority carrier lifetime. The minority carrier lifetime is highly dependent on the material properties of the semiconductor such as doping and surface recombination velocities. In general, the $1/\tau$ term will dominate the denominator of the Δn term until very high frequencies.

From here, we can predict the displacement of the diaphragm similarly to the

thermoelastic effect:

$$w_{PE} = \frac{1}{4} R^2 M_n \Delta n \tag{A.2}$$

where M_n is the moment-generating constant determined from the photoelectric properties of the materials and geometry of the diaphragm. The constant M_n is defined similarly to M_T (described in Section 4.3.3.2):

$$M_n = \frac{\int_0^L d_n(z) E(z - z_n) dz}{\int_0^L E(z)(z - z_n)^2 \, dz} \tag{A.3}$$

where $d_n(z)$ is the coefficient of electronic deformation of the material at each z-coordinate.

As with the photovoltaic effect, the important factor affecting the plasmaelastic effect is the wavelength of incoming light. It is complicated by the fact that longer wavelengths, while generating more minority carriers for the same optical power, will reduce the amount of light absorbed by the diaphragm (I_A). In general, this means that this effect increases at shorter wavelengths, where nearly all the incoming light will be absorbed and generate carriers. Also like the photovoltaic effect, at incoming wavelengths longer than 1100 nm, no carriers can be generated, and this effect drops to zero.

During the investigation, the level of plasmaelastic asymmetry in the microphones was uncertain. In order to determine the contribution of this effect, an experiment was developed to determine if the plasmaelastic contribution was significantly stronger than the thermoelastic contribution. This experiment is described and explained in Section A.2.

A.1.2 Bending Effects in Symmetric Diaphragms

While in Section 4.3.3.2, we discussed the effects of thermoelastic bending due to asymmetric material properties in the diaphragm, there are some bending effects that

can occur due to thermal and charge carrier gradients in a uniform material. These are actually the primary bending terms in most photoacoustic studies looking at thin plates [35], as previous works investigated uniform plates with thicknesses on the order of hundreds of microns. Because of these thicknesses, heat and charge carriers generated near the surface of the plate will take time to diffuse to the rest of the plate. This will lead to stress gradients and bending of the plate, with a surface-averaged quasistatic displacement of [170]:

$$w_{TB} = \frac{3R^2}{L^3}\alpha_T \int_0^L T(z)(z - L/2)dz \tag{A.4}$$

$$w_{PB} = \frac{3R^2}{L^3}d_n \int_0^L \Delta n(z)(z - L/2)dz \tag{A.5}$$

These terms were disregarded due to the thicknesses of the MEMS diaphragms being on the order of microns, and therefore the heat and charge carriers quickly diffuse through the plate. This causes the temperature and carrier concentration to be nearly uniform (i.e. $T(z) = T$ and $\Delta n(z) = \Delta n$), and the bending moment is reduced to zero. If the device in question has significantly higher thicknesses or lower material diffusivity, these bending terms may become a significant contribution to the output signal.

A.1.3 Thermoelastic and Plasmaelastic Expansion

Beyond bending caused by stress gradients, the diaphragm will also displace due to linear expansions and contractions caused by changing temperatures and charge-carrier densities. This was first modeled by [34], showing that displacements will

occur in the z-direction as the material changes size [170]:

$$w_{TX} = \frac{1}{2} L \alpha_T \int_0^L T(z) dz \tag{A.6}$$

$$w_{PX} = \frac{1}{2} L d_n \int_0^L \Delta n(z) dz \tag{A.7}$$

In general, these terms are negligible due to the thinness of the MEMS structures.

The diaphragms will also expand radially, which can also generate a displacement in the z-direction. This effect is difficult to model as it requires specific knowledge of the mechanical boundary conditions and initial curvature of the diaphragm. In general, this is not considered to be a strong contributor to the output signal, as MEMS diaphragms are often designed to be able to expand freely in the radial direction, but this may be a concern in some MEMS devices.

A.1.4 Radiation Pressure

Radiation pressure will affect all the microphones under a laser signal injection attack. Radiation pressure is dependent on the light that is reflected and absorbed by the membrane, as these photons impart momentum into the membrane. Assuming the beam is normal to the plane of the membrane, the equation for the pressure imparted is:

$$P_{RP} = \frac{I_A + 2I_R}{c} \tag{A.8}$$

where c is the speed of light. Note that the first term is due to the absorbed light, and the second term is due to the reflected light. With a quasistatic approximation, the displacement of the membrane due to this pressure is then:

$$w_{RP} = \frac{A_B}{K_d} P_{RP} \tag{A.9}$$

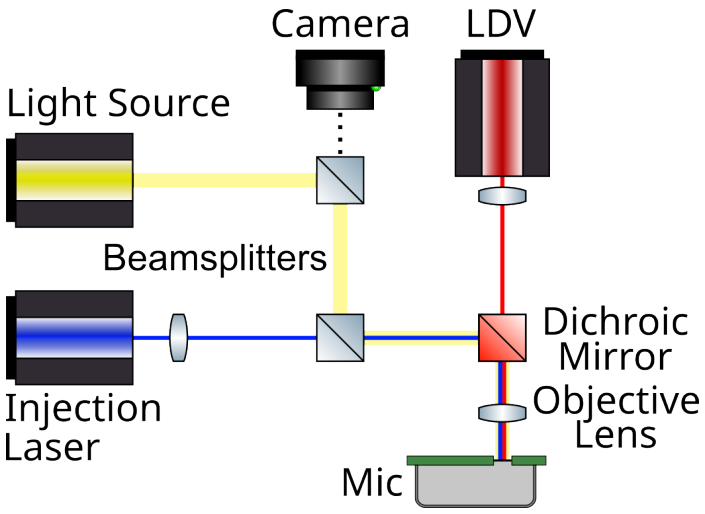

Figure A.1: Optical setup to measure diaphragm displacement while performing laser signal injection.

For normal parameters of MEMS microphones [111], the displacement if the diaphragm would be on the order of a picometer per ten milliwatts of incoming optical power, which is only measurable by extremely sensitive microphones. This means radiation pressure can be safely disregarded in most cases.

A.2 Investigating Plasmaelastic Bending Effects using a Laser Doppler Vibrometer

During the investigation of laser effects into MEMS microphones, plasmaelastic bending of asymmetric diaphragms was considered a viable candidate for the effects seen in laser signal injection. The purpose of this section is to show the investigation that was performed on each of the microphones using a laser doppler vibrometer. Ultimately, there was very little indication of plasmaelastic effects, but the results are presented here.

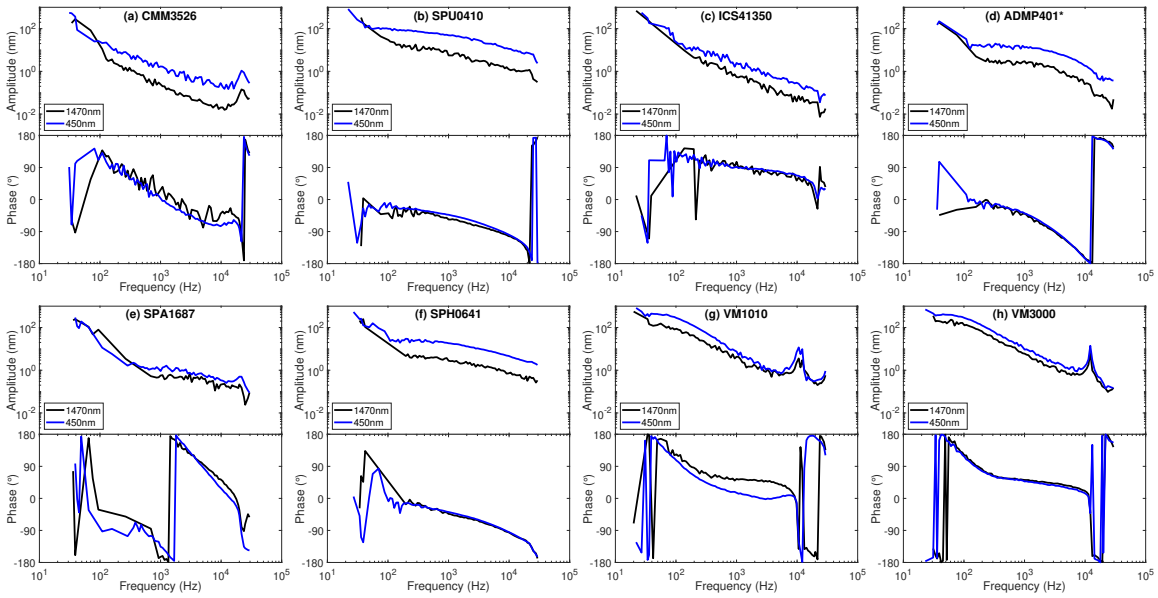

Figure A.2: A comparison of diaphragm displacement with sub-bandgap and super-bandgap lasers. Displacement was measured with a laser doppler vibrometer, which indicated that thermoelastic effects dominated plasmaelastic effects in all microphones. (*) The back package of the ADMP401 was removed to measure the back diaphragm directly.

A.2.1 LDV Setup

To test for plasmaelastic effects as described in Section A.1.1, we used a laser doppler vibrometer to measure the displacement of the diaphragm while under laser signal injection. An overview of the setup is shown in Figure A.1. In the experiment, a Polytec OFV3001 Vibrometer with an OFV303 laser head was used to measure the displacement of the diaphragm. A Thorlabs DMLP550 dichroic mirror with a 550 nm cut-on wavelength was used to combine the measurement beam with the excitation beam. The output of the vibrometer was measured by the Picoscope and captured by the same custom MATLAB program to compare the frequency response of the vibrometer and the microphone output.

A.2.2 Experimental Results

Now that we have a clear idea of the contribution of thermal effects in each of the microphones, the next step in the process is to determine the contribution of plasmaelastic effects.

Plasmaelastic and thermal effects are difficult to separate, as both depend on the amount of light absorbed by the diaphragm (I_A). In order to differentiate between the two, previous works [27, 28] relied on the fact that the coefficient of electronic deformation in silicon is negative, meaning that an increase in the concentration of minority carriers results in a contraction of the material. This contraction will be directly opposed to the expansion caused by the increase in temperature. Because of this, the thermal signal and the plasmaelastic signal can have opposite polarities, which can be used to differentiate the two phenomena. While this is may not always the case depending on the geometry and materials of the diaphragm, a change in polarity would be a clear indication of a significant contribution from the plasmaelastic phenomenon.

We can measure this flip in polarity by using a laser doppler vibrometer (LDV). An LDV uses special optics to determine the velocity of a moving reflective surface. In the case of MEMS microphones, we can measure the velocity of the diaphragm directly as it is excited by the injection laser. By integrating this velocity signal, we can monitor the displacement of the diaphragm (w), ensuring we are isolating photoacoustic effects.

We use the LDV to measure the motion of the diaphragm in two cases. First, we measure the motion of the diaphragm while exciting it with a 1470nm sub-bandgap IR injection laser. This will give us an amplitude and phase response that we can know is due only to thermal effects, as plasmaelastic effects rely on the generation of minority carriers. We repeat the same measurement of the diaphragm using the LDV, only now using a blue 450nm laser with energy above the bandgap, where minority

carriers are generated. If the phase of the signal flips or changes significantly, we can determine that plasmaelastic effects provide a significant contribution towards the laser injection signal.

We performed these two experiments on each of the eight microphones using an LDV to measure the motion of the membrane. The measurement laser was combined with the injection laser using a dichroic mirror, allowing the measurement laser to transmit through the mirror, while the excitation laser is reflected by the mirror. We again used a 5mW output bias with a 1mW signal modulation. A frequency sweep was performed on each of the microphones, and the amplitude and phase response of the diaphragm velocity signal was collected and integrated to determine the diaphragm displacement signal. The Vesper VM1010 and VM3000 were included in the experiment as a control group, as their diaphragms do not contain semiconductor materials, and therefore will not exhibit any plasmaelastic effects.

The results of the experiments are shown in Figure A.2. It can be clearly seen that all the microphones exhibit a measurable motion of the diaphragm. In all cases, the phase response with the 1470nm laser was the same as with the 450nm laser, demonstrating no change in polarity. This indicates that thermal-generated motions of the diaphragms are dominant over plasmaelastic effects in the microphones that we investigated.

Note that there is a significant increase in the amplitude of the displacement signal between the 1470nm and the 450nm experiments. We believe this is due to the fact that significantly more energy will be absorbed by the polysilicon diaphragms at 450nm, as the diaphragms will be partially transparent to the 1470nm laser. Despite the difference in amplitude, the phase responses were consistent between the two injection lasers.

While these results do not completely remove the possibility of plasmaelastic effects, they do provide a contraindication towards plasmaelastic effects being a domi-

nant factor. This supports the theory that thermal effects are the dominant photoacoustic mechanism.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Z. Wang, Q. Zou, Q. Song, and J. Tao, "The era of silicon MEMS microphone and look beyond," in *2015 Transducers - 2015 18th International Conference on Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS)*, June 2015, pp. 375–378.

[2] C. Bolton, K. Fu, J. Hester, and J. Han, "How to curtail oversensing in the home," *Communications of the ACM*, vol. 63, no. 6, pp. 20–24, May 2020. [Online]. Available: https://doi.org/10.1145/3396261

[3] K. Fu and W. Xu, "Risks of trusting the physics of sensors," *Commun. ACM*, vol. 61, no. 2, p. 20–23, Jan. 2018. [Online]. Available: https://doi.org/10.1145/3176402

[4] Y. Cao, C. Xiao, B. Cyr, Y. Zhou, W. Park, S. Rampazzi, Q. A. Chen, K. Fu, and Z. M. Mao, "Adversarial sensor attack on lidar-based perception in autonomous driving," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019, pp. 2267–2281.

[5] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, "Light Commands: Laser-Based Audio Injection Attacks on Voice-Controllable Systems," in *29th USENIX Security Symposium (USENIX Security 20)*, 2020, pp. 2631–2648. [Online]. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/sugawara

[6] B. Cyr, T. Sugawara, and K. Fu, "Why Lasers Inject Perceived Sound Into MEMS Microphones: Indications and Contraindications of Photoacoustic and Photoelectric Effects," in *2021 IEEE Sensors*, Oct. 2021, pp. 1–4, iSSN: 2168-9229.

[7] B. Cyr, Y. Long, T. Sugawara, and K. Fu, "Position Paper: Space System Threat Models Must Account for Satellite Sensor Spoofing," in *SpaceSec23*, Feb. 2023.

[8] C. Yan, H. Shin, C. Bolton, W. Xu, Y. Kim, and K. Fu, "SoK: A Minimalist Approach to Formalizing Analog Sensor Security," in *IEEE S&P 2020*, 2020, pp. 233–248.

[9] A. Einstein, "Zur Quantentheorie der Strahlung. (German) [On the Quantum Theory of Radiation]," *j-PHYSIKAL-Z*, vol. 18, pp. 121–128, 1917.

[10] K. Thyagarajan and A. Ghatak, *Lasers*, ser. Graduate Texts in Physics. Boston, MA: Springer US, 2011. [Online]. Available: http://link.springer.com/10.1007/978-1-4419-6442-7

[11] S. A. Self, "Focusing of spherical Gaussian beams," *Applied Optics*, vol. 22, no. 5, pp. 658–661, Mar. 1983, publisher: Optica Publishing Group. [Online]. Available: https://opg.optica.org/ao/abstract.cfm?uri=ao-22-5-658

[12] L. C. Andrews, W. B. Miller, and J. C. Ricklin, "Geometrical representation of Gaussian beams propagating through complex paraxial optical systems," *Applied Optics*, vol. 32, no. 30, p. 5918, Oct. 1993. [Online]. Available: https://opg.optica.org/abstract.cfm?URI=ao-32-30-5918

[13] P. L. Muiño, "Introducing the Uncertainty Principle Using Diffraction of Light Waves," *Journal of Chemical Education*, vol. 77, no. 8, p. 1025, Aug. 2000. [Online]. Available: https://pubs.acs.org/doi/abs/10.1021/ed077p1025

[14] J. Peatross and M. Ware, *Physics of Light and Optics*, 2013.

[15] T. G. Mayerhöfer, S. Pahlow, and J. Popp, "The Bouguer-Beer-Lambert Law: Shining Light on the Obscure," *ChemPhysChem*, vol. 21, no. 18, pp. 2029–2046, Sep. 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/cphc.202000464

[16] H. Hertz, "Ueber einen Einfluss des ultravioletten Lichtes auf die electrische Entladung," *Annalen der Physik*, vol. 267, no. 8, pp. 983–1000, 1887, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.18872670827. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.18872670827

[17] A. Einstein, "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt," *Annalen der Physik*, vol. 322, no. 6, pp. 132–148, 1905, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19053220607. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19053220607

[18] M. Planck, "Ueber das Gesetz der Energieverteilung im Normalspectrum," *Annalen der Physik*, vol. 309, no. 3, pp. 553–563, 1901, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19013090310. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19013090310

[19] C. Honsberg and S. Bowden, "PVEducation," 2019. [Online]. Available: https://www.pveducation.org/

[20] W. Smith, "Effect of Light on Selenium During the Passage of An Electric Current *," *Nature*, vol. 7, no. 173, pp. 303–303, Feb. 1873,

number: 173 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/007303e0

[21] S. O. Kasap, "Photoconductivity: Fundamental Concepts," in *Photoconductivity and Photoconductive Materials*. John Wiley & Sons, Ltd, 2022, pp. 1–88, section: 1 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119579182.ch1. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119579182.ch1

[22] W. B. Gauster and D. H. Habing, "Electronic Volume Effect in Silicon," *Physical Review Letters*, vol. 18, no. 24, pp. 1058–1061, Jun. 1967, publisher: American Physical Society. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.18.1058

[23] V. E. Gusev and A. A. Karabutov, "Laser optoacoustics," *NASA STI/Recon Technical Report A*, vol. 93, Jul. 1991. [Online]. Available: http://adsabs.harvard.edu/abs/1991STIA...9316842G

[24] D. M. Todorovic-acute, P. M. Nikolic-acute, A. I. Bojicic-acute, and K. T. Radulovic-acute, "Thermoelastic and electronic strain contributions to the frequency transmission photoacoustic effect in semiconductors," *Physical Review B*, vol. 55, no. 23, pp. 15 631–15 642, Jun. 1997, publisher: American Physical Society. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevB.55.15631

[25] D. M. Todorović, M. D. Rabasović, and D. D. Markushev, "Photoacoustic elastic bending in thin film—Substrate system," *Journal of Applied Physics*, vol. 114, no. 21, p. 213510, Dec. 2013, publisher: American Institute of Physics. [Online]. Available: https://aip.scitation.org/doi/10.1063/1.4839835

[26] D. M. Todorovic, M. D. Rabasovic, D. D. Markushev, V. Jovic, K. T. Radulovic, and M. Sarajlic, "Photoacoustic Elastic Bending Method: Characterization of Thin Films on Silicon Membranes," *International Journal of Thermophysics*, vol. 36, no. 5, pp. 1016–1028, Jun. 2015. [Online]. Available: https://doi.org/10.1007/s10765-014-1801-3

[27] P. G. Datskos, S. Rajic, and I. Datskou, "Photoinduced and thermal stress in silicon microcantilevers," *Applied Physics Letters*, vol. 73, no. 16, pp. 2319–2321, Oct. 1998. [Online]. Available: http://aip.scitation.org/doi/10.1063/1.121809

[28] V. Chenniappan, G. A. Umana-Membreno, K. K. M. B. Dilusha Silva, H. Kala, A. J. Keating, M. Martyniuk, J. M. Dell, and L. Faraone, "Characterization and Modeling of Photostriction in Silicon Cantilevers Fabricated on Silicon-on-Insulator Substrates," *Journal of Microelectromechanical Systems*, vol. 24, no. 1, pp. 182–191, Feb. 2015, conference Name: Journal of Microelectromechanical Systems.

[29] M. Faraday, "XX. Experimental researches in electricity.—fourth series," *Philosophical Transactions of the Royal Society of London*, vol. 123, pp. 507–522, 1833, _eprint: https://royalsocietypublishing.org/doi/pdf/10.1098/rstl.1833.0022. [Online]. Available: https://royalsocietypublishing.org/doi/abs/10.1098/rstl.1833.0022

[30] J. S. Steinhart and S. R. Hart, "Calibration curves for thermistors," *Deep Sea Research and Oceanographic Abstracts*, vol. 15, no. 4, pp. 497–503, Aug. 1968. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0011747168900570

[31] N. Noda, R. B. Hetnarski, and Y. Tanigawa, *Thermal stresses*, 2nd ed. New York: Taylor & Francis, 2003.

[32] S. Dutta and A. Pandey, "Overview of residual stress in MEMS structures: Its origin, measurement, and control," *Journal of Materials Science: Materials in Electronics*, vol. 32, no. 6, pp. 6705–6741, Mar. 2021. [Online]. Available: https://doi.org/10.1007/s10854-021-05405-8

[33] R. M. White, "Generation of Elastic Waves by Transient Surface Heating," *Journal of Applied Physics*, vol. 34, no. 12, pp. 3559–3567, Dec. 1963. [Online]. Available: http://aip.scitation.org/doi/10.1063/1.1729258

[34] F. A. McDonald and G. C. Wetsel, "Generalized theory of the photoacoustic effect," *Journal of Applied Physics*, vol. 49, no. 4, pp. 2313–2322, Apr. 1978, publisher: American Institute of Physics. [Online]. Available: https://aip.scitation.org/doi/10.1063/1.325116

[35] G. Rousset, F. Lepoutre, and L. Bertrand, "Influence of thermoelastic bending on photoacoustic experiments related to measurements of thermal diffusivity of metals," *Journal of Applied Physics*, vol. 54, no. 5, pp. 2383–2391, May 1983, publisher: American Institute of Physics. [Online]. Available: https://aip.scitation.org/doi/10.1063/1.332352

[36] S. Manohar and D. Razansky, "Photoacoustics: a historical review," *Advances in Optics and Photonics*, vol. 8, no. 4, p. 586, Dec. 2016. [Online]. Available: https://www.osapublishing.org/abstract.cfm?URI=aop-8-4-586

[37] A. G. Bell, "Upon the production and reproduction of sound by light," *Journal of the Society of Telegraph Engineers*, vol. 9, no. 34, pp. 404–426, 1880.

[38] A. Rosencwaig and A. Gersho, "Theory of the photoacoustic effect with solids," *J. Appl. Phys.*, vol. 47, p. 7, 1976.

[39] G. A. Askar'yan and E. M. Moroz, "Pressure on Evaporation of Matter in a Radiation Beam," *Soviet Journal of Experimental and Theoretical Physics*, vol. 16, p. 1638, Jan. 1963, aDS Bibcode: 1963JETP...16.1638A. [Online]. Available: https://ui.adsabs.harvard.edu/abs/1963JETP...16.1638A

[40] F. Brech and L. Cross, "Optical Microemission Stimulated by a Ruby Laser," in *Applied Spectroscopy*, 1962, vol. 16, no. 2, p. 59.

[41] J. Ready, *Effects of High-Power Laser Radiation*. Saint Louis: Elsevier Science, Aug. 1971, oCLC: 1044729011.

[42] D. E. Fratanduono, T. R. Boehly, P. M. Celliers, M. A. Barrios, J. H. Eggert, R. F. Smith, D. G. Hicks, G. W. Collins, and D. D. Meyerhofer, "The direct measurement of ablation pressure driven by 351-nm laser radiation," *Journal of Applied Physics*, vol. 110, no. 7, p. 073110, Oct. 2011, publisher: American Institute of Physics. [Online]. Available: https://aip.scitation.org/doi/10.1063/1.3646554

[43] A. Einstein, "Zur Quantentheorie der Strahlung. (German) [On the Quantum Theory of Radiation]," *j-PHYSIKAL-Z*, vol. 18, pp. 121–128, 1917.

[44] F. Beaudoin and E. Cole, "Physics of Laser-Based Failure Analysis," in *Microelectronics Failure Analysis*, 7th ed., T. Gandhi, Ed. ASM International, 2019, pp. 196–208. [Online]. Available: http://dl.asminternational.org/handbooks/book/87/chapter/2048400/Physics-of-LaserBased-Failure-Analysis

[45] D. H. Habing, "The use of lasers to simulate radiation-induced transients in semiconductor devices and circuits," *IEEE Transactions on Nuclear Science*, vol. 12, no. 5, pp. 91–100, 1965.

[46] S. P. Skorobogatov and R. J. Anderson, "Optical fault induction attacks," in *Cryptographic Hardware and Embedded Systems - CHES 2002, 4th International Workshop, Redwood Shores, CA, USA, August 13-15, 2002, Revised Papers*, 2002, pp. 2–12.

[47] M. Agoyan, J.-M. Dutertre, A.-P. Mirbaha, D. Naccache, A.-L. Ribotta, and A. Tria, "Single-bit DFA using multiple-byte laser fault injection," in *2010 IEEE International Conference on Technologies for Homeland Security (HST)*, Nov. 2010, pp. 113–119.

[48] V. Beroulle, P. Candelier, S. De Castro, G. Di Natale, J.-M. Dutertre, M.-L. Flottes, D. Hély, G. Hubert, R. Leveugle, F. Lu, P. Maistri, A. Papadimitriou, B. Rouzeyre, C. Tavernier, and P. Vanhauwaert, "Laser-Induced Fault Effects in Security-Dedicated Circuits," in *VLSI-SoC: Internet of Things Foundations*, ser. IFIP Advances in Information and Communication Technology, L. Claesen, M.-T. Sanz-Pascual, R. Reis, and A. Sarmiento-Reyes, Eds. Cham: Springer International Publishing, 2015, pp. 220–240.

[49] K. Yamashita, B. Cyr, K. Fu, W. Burleson, and T. Sugawara, "Redshift: Manipulating Signal Propagation Delay via Continuous-Wave Lasers," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 463–489, Aug. 2022. [Online]. Available: https://tches.iacr.org/index.php/TCHES/article/view/9828

[50] D. Karaklajić, J. Schmidt, and I. Verbauwhede, "Hardware designer's guide to fault attacks," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 21, no. 12, pp. 2295–2306, 2013.

[51] J.-M. Dutertre, J. J. Fournier, A.-P. Mirbaha, D. Naccache, J.-B. Rigaud, B. Robisson, and A. Tria, "Review of fault injection mechanisms and consequences on countermeasures design," in *2011 6th International Conference on Design & Technology of Integrated Systems in Nanoscale Era (DTIS)*. IEEE, 2011, pp. 1–6.

[52] C. Yan, H. Shin, C. Bolton, W. Xu, Y. Kim, and K. Fu, "SoK: A Minimalist Approach to Formalizing Analog Sensor Security," in *2020 IEEE Symposium on Security and Privacy (SP)*, May 2020, pp. 233–248, iSSN: 2375-1207.

[53] D. F. Kune, J. Backes, S. S. Clark, D. Kramer, M. Reynolds, K. Fu, Y. Kim, and W. Xu, "Ghost talk: Mitigating emi signal injection attacks against analog sensors," in *Security and Privacy (SP), 2013 IEEE Symposium on*. IEEE, 2013, pp. 145–159.

[54] C. Kasmi and J. Lopes Esteves, "IEMI Threats for Information Security: Remote Command Injection on Modern Smartphones," *IEEE Transactions on Electromagnetic Compatibility*, vol. 57, no. 6, pp. 1752–1755, Dec. 2015, conference Name: IEEE Transactions on Electromagnetic Compatibility.

[55] R. Chauhan, *A platform for false data injection in frequency modulated continuous wave radar*. Utah State University, 2014.

[56] C. Yan, W. Xu, and J. Liu, "Can You Trust Autonomous Vehicles : Contactless Attacks against Sensors of Self-driving Vehicle," in *DEF CON 24*, 2016.

[57] Y. Shoukry, P. Martin, P. Tabuada, and M. Srivastava, "Non-invasive spoofing attacks for anti-lock braking systems," in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2013, pp. 55–72.

[58] J. Selvaraj, G. Y. Dayanıklı, N. P. Gaunkar, D. Ware, R. M. Gerdes, and M. Mina, "Electromagnetic Induction Attacks Against Embedded Systems," in *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, ser. ASIACCS '18. Incheon, Republic of Korea: Association for Computing Machinery, May 2018, pp. 499–510. [Online]. Available: https://doi.org/10.1145/3196494.3196556

[59] S. Maruyama, S. Wakabayashi, and T. Mori, "Tap 'n Ghost: A Compilation of Novel Attack Techniques against Smartphone Touchscreens," in *2019 IEEE Symposium on Security and Privacy (SP)*. San Francisco, CA, USA: IEEE, May 2019, pp. 620–637. [Online]. Available: https://ieeexplore.ieee.org/document/8835251/

[60] H. Shan, B. Zhang, Z. Zhan, D. Sullivan, S. Wang, and Y. Jin, "Invisible Finger: Practical Electromagnetic Interference Attack on Touchscreen-based Electronic Devices," in *2022 IEEE Symposium on Security and Privacy (SP)*. San Francisco, CA, USA: IEEE, May 2022, pp. 1246–1262. [Online]. Available: https://ieeexplore.ieee.org/document/9833718/

[61] K. Wang, R. Mitev, C. Yan, X. Ji, A.-R. Sadeghi, and W. Xu, "GhostTouch: Targeted Attacks on Touchscreens without Physical Touch," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 1543–1559. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/wang-kai

[62] Y. Tu, S. Rampazzi, B. Hao, A. Rodriguez, K. Fu, and X. Hei, "Trick or Heat? Manipulating Critical Temperature-based Control Systems Using Rectification Attacks," in *CCS 2019*, 2019, pp. 2301–2315.

[63] Y. Son, H. Shin, D. Kim, Y.-S. Park, J. Noh, K. Choi, J. Choi, Y. Kim *et al.*, "Rocking drones with intentional sound noise on gyroscopic sensors." in *USENIX Security Symposium*, 2015, pp. 881–896.

[64] T. Trippel, O. Weisse, W. Xu, P. Honeyman, and K. Fu, "Walnut: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks," in *Security and Privacy (EuroS&P), 2017 IEEE European Symposium on*. IEEE, 2017, pp. 3–18.

[65] Y. Tu, Z. Lin, I. Lee, and X. Hei, "Injected and Delivered: Fabricating Implicit Control over Actuation Systems by Spoofing Inertial Sensors," in *27th USENIX Security Symposium*, 2018, pp. 1545–1562. [Online]. Available: https://www.usenix.org/conference/usenixsecurity18/presentation/tu

[66] Z. Wang, K. Wang, B. Yang, S. Li, and A. Pan, "Sonic Gun to Smart Devices: Your Devices Lose Control Under Ultrasound/Sound by SonicGun," in *Black Hat, USA*, 2017. [Online]. Available: https://sonicgun.github.io/

[67] S. Nashimoto, D. Suzuki, T. Sugawara, and K. Sakiyama, "Sensor CON-Fusion: Defeating kalman filter in signal injection attack," in *Proceedings of the 2018 on Asia Conference on Computer and Communications Security, AsiaCCS 2018*, 2018, pp. 511–524.

[68] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, "Dolphinattack: Inaudible Voice Commands," in *CCS 2017*, 2017, pp. 103–117.

[69] L. Song and P. Mittal, "Inaudible Voice Commands," *arXiv:1708.07238 [cs]*, Aug. 2017, arXiv: 1708.07238. [Online]. Available: http://arxiv.org/abs/1708.07238

[70] N. Roy, S. Shen, H. Hassanieh, and R. R. Choudhury, "Inaudible Voice Commands: The Long-Range Attack and Defense," in *15th USENIX*

Symposium on Networked Systems Design and Implementation (NSDI 18), 2018, pp. 547–560. [Online]. Available: https://www.usenix.org/conference/nsdi18/presentation/roy

[71] C. Yan, G. Zhang, X. Ji, T. Zhang, T. Zhang, and W. Xu, "The Feasibility of Injecting Inaudible Voice Commands to Voice Assistants," *IEEE Transactions on Dependable and Secure Computing*, pp. 1–1, 2019, conference Name: IEEE Transactions on Dependable and Secure Computing.

[72] W. Xu, C. Yan, W. Jia, X. Ji, and J. Liu, "Analyzing and Enhancing the Security of Ultrasonic Sensors for Autonomous Vehicles," *IEEE Internet of Things Journal*, vol. 5, no. 6, pp. 5015–5029, Dec. 2018, conference Name: IEEE Internet of Things Journal.

[73] C. Bolton, S. Rampazzi, C. Li, A. Kwong, W. Xu, and K. Fu, "Blue note: How intentional acoustic interference damages availability and integrity in hard disk drives and operating systems," in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2018, pp. 1048–1062.

[74] X. Ji, Y. Cheng, Y. Zhang, K. Wang, C. Yan, W. Xu, and K. Fu, "Poltergeist: Acoustic Adversarial Machine Learning against Cameras and Computer Vision," in *2021 IEEE Symposium on Security and Privacy (SP)*, May 2021, pp. 160–175, iSSN: 2375-1207.

[75] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote attacks on automated vehicles sensors: Experiments on camera and LiDAR," *Black Hat Europe*, vol. 11, p. 2015, 2015.

[76] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," *DEFCON*, vol. 24, 2016.

[77] B. Nassi, D. Nassi, R. Ben-Netanel, Y. Mirsky, O. Drokin, and Y. Elovici, "Phantom of the ADAS: Phantom Attacks on Driver-Assistance Systems," *IACR Cryptology ePrint Archive*, 2020.

[78] Y. Man, M. Li, and R. Gerdes, "GhostImage: Remote Perception Attacks against Camera-based Image Classification Systems," in *RAID 2020*, 2020, pp. 317–332.

[79] S. Köhler, G. Lovisotto, S. Birnbach, R. Baker, and I. Martinovic, "They See Me Rollin': Inherent Vulnerability of the Rolling Shutter in CMOS Image Sensors," Dec. 2021, arXiv:2101.10011 [cs]. [Online]. Available: http://arxiv.org/abs/2101.10011

[80] C. Yan, Z. Xu, Z. Yin, X. Ji, and W. Xu, "Rolling Colors: Adversarial Laser Exploits against Traffic Light Recognition," in *USENIX Security 22*, Aug. 2022, pp. 1957–1974.

[81] T. Sato, S. H. V. Bhupathiraju, M. Clifford, T. Sugawara, Q. A. Chen, and S. Rampazzi, "WIP: Infrared Laser Reflection Attack Against Traffic Sign Recognition Systems," in *Proceedings Inaugural International Symposium on Vehicle Security & Privacy*. San Diego, CA, USA: Internet Society, 2023. [Online]. Available: https://www.ndss-symposium.org/wp-content/uploads/2023/02/vehiclesec2023-23055-paper.pdf

[82] Y.-S. Park, Y. Son, H. Shin, D. Kim, and Y. Kim, "This ain't your dose: Sensor spoofing attack on medical infusion pump." in *WOOT*, 2016.

[83] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and Dazzle: Adversarial Optical Channel Exploits Against Lidars for Automotive Applications," in *CHES 2017*, 2017, pp. 445–467.

[84] H. Shin, J. Noh, D. Kim, and Y. Kim, "The System That Cried Wolf: Sensor Security Analysis of Wide-area Smoke Detectors for Critical Infrastructure," 2020, report Number: 562. [Online]. Available: https://eprint.iacr.org/2020/562

[85] T. Tanaka and T. Sugawara, "Laser-Based Signal-Injection Attack on Piezoresistive MEMS Pressure Sensors," in *2022 IEEE Sensors*, Oct. 2022, pp. 1–4.

[86] Y. Shoukry, P. Martin, Y. Yona, S. Diggavi, and M. Srivastava, "Pycra: Physical challenge-response authentication for active sensors under spoofing attacks," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 1004–1015.

[87] I. Giechaskiel, Y. Zhang, and K. B. Rasmussen, "A framework for evaluating security in the presence of signal injection attacks," *CoRR*, vol. abs/1901.03675, 2019.

[88] J. Sun, Y. Cao, Q. A. Chen, and Z. M. Mao, "Towards Robust {LiDAR-based} Perception in Autonomous Driving: General Black-box Adversarial Sensor Attack and Countermeasures," in *29th USENIX Security Symposium*, 2020, pp. 877–894. [Online]. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/sun

[89] H. Choi, W.-C. Lee, Y. Aafer, F. Fei, Z. Tu, X. Zhang, D. Xu, and X. Deng, "Detecting Attacks Against Robotic Vehicles: A Control Invariant Approach," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. Toronto Canada: ACM, Oct. 2018, pp. 801–816. [Online]. Available: https://dl.acm.org/doi/10.1145/3243734.3243752

[90] P. Guo, H. Kim, N. Virani, J. Xu, M. Zhu, and P. Liu, "RoboADS: Anomaly Detection Against Sensor and Actuator Misbehaviors in Mobile Robots," in *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Jun. 2018, pp. 574–585, iSSN: 2158-3927.

[91] H. Choi, S. Kate, Y. Aafer, X. Zhang, and D. Xu, "Software-based Realtime Recovery from Sensor Attacks on Robotic Vehicles," in *23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*, 2020, pp. 349–364. [Online]. Available: https://www.usenix.org/conference/raid2020/presentation/choi

[92] "Baidu starts mass production of autonomous buses," https://www.dw.com/en/baidu-starts-mass-production-of-autonomous-buses/a-44525629, 2018.

[93] "Waymo's autonomous cars have driven 8 million miles on public roads," https://www.theverge.com/2018/7/20/17595968/waymo-self-driving-cars-8-million-miles-testing, 2018.

[94] "You can take a ride in a self-driving Lyft during CES," https://www.theverge.com/2018/1/2/16841090/lyft-aptiv-self-driving-car-ces-2018, 2018.

[95] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications," in *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, 2017, pp. 445–467.

[96] X. Zhu and S. Xiang, "Adaptive coding for lidar systems," US Patent US10 466 342B1, Nov., 2019. [Online]. Available: https://patents.google.com/patent/US10466342B1/en?assignee=hesai&oq=hesai

[97] Y. Cao, S. H. Bhupathiraju, P. Naghavi, T. Sugawara, Z. M. Mao, and S. Rampazzi, "You Can't See Me: Physical Removal Attacks on LiDAR-based Autonomous Vehicles Driving Frameworks," Oct. 2022, arXiv:2210.09482 [cs] version: 1. [Online]. Available: http://arxiv.org/abs/2210.09482

[98] W. Diao, X. Liu, Z. Zhou, and K. Zhang, "Your voice assistant is mine: How to abuse speakers to steal information and control your phone," in *Proceedings of the 4th ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*. ACM, 2014, pp. 63–74.

[99] Y. Jang, C. Song, S. P. Chung, T. Wang, and W. Lee, "A11y attacks: Exploiting accessibility in operating systems," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 103–115.

[100] N. Carlini, P. Mishra, T. Vaidya, Y. Zhang, M. Sherr, C. Shields, D. Wagner, and W. Zhou, "Hidden voice commands." in *USENIX Security Symposium*, 2016, pp. 513–530.

[101] T. Vaidya, Y. Zhang, M. Sherr, and C. Shields, "Cocaine noodles: exploiting the gap between human and machine speech recognition," *Presented at WOOT*, vol. 15, pp. 10–11, 2015.

[102] X. Yuan, Y. Chen, Y. Zhao, Y. Long, X. Liu, K. Chen, S. Zhang, H. Huang, X. Wang, and C. A. Gunter, "CommanderSong: A systematic approach for practical adversarial voice recognition," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 49–64.

[103] D. Kumar, R. Paccagnella, P. Murley, E. Hennenfent, J. Mason, A. Bates, and M. Bailey, "Skill squatting attacks on Amazon Alexa," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 33–47.

[104] N. Zhang, X. Mi, X. Feng, X. Wang, Y. Tian, and F. Qian, "Understanding and mitigating the security risks of voice-controlled third-party skills on amazon alexa and google home," *arXiv preprint arXiv:1805.01525*, 2018.

[105] N. Roy, H. Hassanieh, and R. Roy Choudhury, "Backdoor: Making microphones hear inaudible sounds," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 2017, pp. 2–14.

[106] L. Song and P. Mittal, "Inaudible voice commands," *arXiv preprint arXiv:1708.07238*, 2017.

[107] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, "DolphinAttack: Inaudible voice commands," in *ACM Conference on Computer and Communications Security*. ACM, 2017, pp. 103–117.

[108] N. Roy, S. Shen, H. Hassanieh, and R. R. Choudhury, "Inaudible voice commands: The long-range attack and defense," in *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, 2018, pp. 547–560.

[109] C. Kasmi and J. L. Esteves, "Iemi threats for information security: Remote command injection on modern smartphones," *IEEE Transactions on Electromagnetic Compatibility*, vol. 57, no. 6, pp. 1752–1755, 2015.

[110] Y. Wang, H. Guo, and Q. Yan, "GhostTalk: Interactive Attack on Smartphone Voice System Through Power Line," *Network and Distributed Systems Security Symposium 2022*, p. 15, Apr. 2022.

[111] N. N. Peña-García, L. A. Aguilera-Cortés, M. A. González-Palacios, J.-P. Raskin, and A. L. Herrera-May, "Design and Modeling of a MEMS Dual-Backplate Capacitive Microphone with Spring-Supported Diaphragm for Mobile Device Applications," *Sensors*, vol. 18, no. 10, p. 3545, Oct. 2018.

[112] S. Shubham, Y. Seo, V. Naderyan, X. Song, A. J. Frank, J. T. M. G. Johnson, M. da Silva, and M. Pedersen, "A Novel MEMS Capacitive Microphone with Semiconstrained Diaphragm Supported with Center and Peripheral Backplate Protrusions," *Micromachines*, vol. 13, no. 1, p. 22, Jan. 2022.

[113] P. Loeppert and S. Lee, "SIS0NIC - THE FIRST COMMERCIALIZED MEMS MICROPHONE," in *2006 Solid-State, Actuators, and Microsystems Workshop*

Technical Digest. Hilton Head, South Carolina, USA: Transducer Research Foundation, Inc., Jun. 2006, pp. 27–30.

[114] J. Weigold, T. Brosnihan, J. Bergeron, and X. Zhang, "A MEMS Condenser Microphone for Consumer Applications," in *19th IEEE International Conference on Micro Electro Mechanical Systems*. Istanbul, Turkey: IEEE, 2006, pp. 86–89.

[115] K. Grosh and R. J. Littrell, "Piezoelectric MEMS microphone," US Patent US10 964 880B2, Mar., 2021.

[116] L. Zhang, S. Tan, J. Yang, and Y. Chen, "Voicelive: A phoneme localization based liveness detection for voice authentication on smartphones," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1080–1091.

[117] L. Zhang, S. Tan, and J. Yang, "Hearing your voice is not enough: An articulatory gesture based liveness detection for voice authentication," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*. ACM, 2017, pp. 57–71.

[118] L. Lu, J. Yu, Y. Chen, H. Liu, Y. Zhu, Y. Liu, and M. Li, "Lippass: Lip reading-based user authentication on smartphones leveraging acoustic signals," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 1466–1474.

[119] L. Lu, J. Yu, Y. Chen, H. Liu, Y. Zhu, L. Kong, and M. Li, "Lip reading-based user authentication through acoustic sensing on smartphones," *IEEE/ACM Transactions on Networking*, vol. 27, no. 1, pp. 447–460, 2019.

[120] H. Feng, K. Fawaz, and K. Shin, "Continuous authentication for voice assistants," in *ACM MobiCom*, 2017.

[121] D. Davidson, H. Wu, R. Jellinek, T. Ristenpart, and V. Singh, "Controlling UAVs with sensor input spoofing attacks," in *Proceedings of the 10th USENIX Conference on Offensive Technologies*, ser. WOOT'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 221–231. [Online]. Available: http://dl.acm.org/citation.cfm?id=3027019.3027039

[122] M. Scholl, "Introduction to Cybersecurity for Commercial Satellite Operations," NIST, Tech. Rep., 2021.

[123] "Nanosats Database." [Online]. Available: https://www.nanosats.eu/

[124] M. Manulis, C. P. Bridges, R. Harrison, V. Sekar, and A. Davis, "Cyber Security in New Space," *International Journal of Information Security*, vol. 20, no. 3, pp. 287–311, 2021.

[125] J. Pavur and I. Martinovic, "Building a Launchpad for Satellite Cyber-Security Research: Lessons from 60 Years of Spaceflight," *Journal of Cybersecurity*, vol. 8, no. 1, 2022.

[126] P. Tedeschi, S. Sciancalepore, and R. Di Pietro, "Satellite-based Communications Security: A Survey of Threats, Solutions, and Research Challenges," *Computer Networks*, p. 109246, 2022.

[127] S. Theil, P. Appel, and A. Schleicher, "Low Cost, Good Accuracy - Attitude Determination using Magnetometer and Simple Sun Sensor," in *17th Annual AIAA/USU Conference on Small Satellites*, 2003, p. 10.

[128] J. H. Wessels, "Infrared horizon sensor for CubeSat implementation," Thesis, Stellenbosch : Stellenbosch University, Mar. 2018, accepted: 2018-02-20T09:35:09Z. [Online]. Available: https://scholar.sun.ac.za:443/handle/10019.1/103564

[129] C. R. McBryde and E. G. Lightsey, "A star tracker design for CubeSats," in *2012 IEEE Aerospace Conference*, Mar. 2012, pp. 1–14, iSSN: 1095-323X.

[130] M. Zahran and M. Aly, "A Solar Cell Based Coarse Sun Sensor for a Small LEO Satellite Attitude Determination," *Journal of Power Electronics*, vol. 9, no. 4, p. 12, 2009.

[131] H. Kaushal and G. Kaddoum, "Optical Communication in Space: Challenges and Mitigation Techniques," *IEEE Communications Surveys Tutorials*, vol. 19, no. 1, pp. 57–96, 2017, conference Name: IEEE Communications Surveys Tutorials.

[132] D. G. Aviv, *Laser space communications*, ser. Artech House space technology and applications series. Boston: Artech House, 2006, oCLC: ocm70403107.

[133] S. Qian, "Hyperspectral Satellites, Evolution, and Development History," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 7032–7056, 2021.

[134] A. J. Bell, "Analysis of GPS Satellite Allocation for the United States Nuclear Detonation Detection System (USNDS)," Air Force Institute of Technology, Tech. Rep., Mar. 2002. [Online]. Available: https://apps.dtic.mil/sti/citations/ADA401805

[135] J. Menn, "China-based campaign breached satellite, defense companies: Symantec," *Reuters*, Jun. 2018. [Online]. Available: https://www.reuters.com/article/us-china-usa-cyber-idUSKBN1JF2X0

[136] D. Bird, "Cybersecurity Considerations for Internet of Things Small Satellite Systems," *Current Analysis on Communication Engineering*, p. 11, Mar. 2019.

[137] S. Visner and S. C. K. Ph.D, "Cyber Best Practices for Small Satellites," Nov. 2020. [Online]. Available: https://www.mitre.org/publications/technical-papers/cyber-best-practices-for-small-satellites

[138] E. J. Birrane, S. Heiner, J. Hopkins, and E. Birrane, "A Novel Approach to Transport-Layer Security for Spacecraft Constellations," *34th Annual Small Satellite Conference*, p. 14, 2020.

[139] S. Kinser, P. de Graaf, M. Stein, F. Hughey, R. Roller, D. Voss, and A. Salmoiraghi, "Scoring Trust Across Hybrid-Space: A Quantitative Framework Designed to Calculate Cybersecurity Ratings, Measures, and Metrics to Inform a Trust Score," *34th Annual Small Satellite Conference*, p. 8, 2020.

[140] F. Schubert, "Satellite cyber security is more important than ever – here's why," Jan. 2022. [Online]. Available: https://airbus-cyber-security.com/satellite-cyber-security-is-more-important-than-ever-heres-why/

[141] J. A. Steinberger, "A Survey of Satellite Communications System Vulnerabilities," Ph.D. dissertation, Air Force Institute of Technology, Jun. 2008.

[142] J. R. Solin, "Ground Based Laser Triggered Discharges on Satellite Solar Arrays," in *Laser-Induced Damage in Optical Materials: 2005*, vol. 5991. SPIE, 2006, pp. 723–731.

[143] J. Solin, "Airborne Laser Threat to Commercial Space Telescopes," *Optical Engineering*, vol. 53, no. 9, p. 095105, 2014.

[144] S. Wang and L. Guo, "Analysis of Laser Jamming to Satellite-based Detector," in *International Symposium on Photoelectronic Detection and Imaging 2009*, vol. 7382. SPIE, 2009, pp. 805–813.

[145] Z. Liu, C. Lin, and G. Chen, "Space Attack Technology Overview," in *Journal of Physics*, vol. 1544, no. 1, 2020, p. 012178.

[146] D. H. Hilland, G. S. Phipps, C. M. Jingle, and G. Newton, "Satellite Threat Warning and Attack Reporting," in *1998 IEEE Aerospace Conference Proceedings*, vol. 2, 1998, pp. 207–217.

[147] J. Altmann, "Offensive Capabilities of Space-Based Lasers," *Bulletin of Peace Proposals*, vol. 17, no. 2, pp. 151–158, 1986.

[148] Y. Park, Y. Son, H. Shin, D. Kim, and Y. Kim, "This Ain't Your Dose: Sensor Spoofing Attack on Medical Infusion Pump," in *WOOT 2016*, 2016.

[149] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, "Light Commands: Laser-Based Audio Injection Attacks on Voice-Controllable Systems," in *USENIX Security 20*, Aug. 2020, pp. 2631–2648.

[150] Y. Hayashi, N. Homma, T. Mizuki, T. Aoki, and H. Sone, "Transient IEMI Threats for Cryptographic Devices," *IEEE transactions on Electromagnetic Compatibility*, vol. 55, no. 1, pp. 140–148, 2012.

[151] Y. Long, S. Rampazzi, T. Sugawara, and K. Fu, "Protecting COVID-19 Vaccine Transportation and Storage from Analog Cybersecurity Threats," *Biomedical Instrumentation & Technology*, vol. 55, no. 3, pp. 112–117, 2021.

[152] Q. Jiang, X. Ji, C. Yan, Z. Xie, H. Lou, and W. Xu, "GlitchHiker: Uncovering Vulnerabilities of Image Signal Transmission with IEMI," in *USENIX Security 23*, 2023.

[153] P. Ortega, G. López-Rodríguez, J. Ricart, M. Domínguez, L. M. Castañer, J. M. Quero, C. L. Tarrida, J. García, M. Reina, A. Gras, and M. Angulo, "A Miniaturized Two Axis Sun Sensor for Attitude Control of Nano-Satellites," *IEEE Sensors Journal*, vol. 10, no. 10, pp. 1623–1632, Oct. 2010.

[154] A. N. A. Ali, M. H. Saied, M. Z. Mostafa, and T. M. Abdel-Moneim, "A Survey of Maximum PPT Techniques of PV Systems," in *2012 IEEE Energytech*, May 2012, pp. 1–17.

[155] J. Cao, J. Yang, S. Yuan, X. Shen, Y. Liu, C. Yan, W. Li, and T. Chen, "In-Flight Observations of Electromagnetic Interferences Emitted by Satellite," *Science in China Series E: Technological Sciences*, vol. 52, no. 7, pp. 2112–2118, Jul. 2009.

[156] A. Keshmirian, "A Physically-based Approach for Lens Flare Simulation," Ph.D. dissertation, UC San Diego, 2008.

[157] A. K. Maini, *Handbook of Defence Electronics and Optronics: Fundamentals, Technologies and Systems*. Wiley, Apr. 2018.

[158] M. Wacks, "The Alpha Program," *Journal of Directed Energy*, 2006.

[159] B. Hendrickx, "The Space Review: Peresvet: a Russian Mobile Laser System to Dazzle Enemy Satellites," Jun. 2020. [Online]. Available: https://www.thespacereview.com/article/3967/1

[160] "NRO Confirms Chinese Laser Test Illuminated U.S. Spacecraft," Oct. 2006. [Online]. Available: https://spacenews.com/nro-confirms-chinese-laser-test-illuminated-us-spacecraft/

[161] M. N. Zervas and C. A. Codemard, "High Power Fiber Lasers: A Review," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 20, no. 5, pp. 219–241, Sep. 2014.

[162] K. D. Skelley, "Directed Energy Weapon System Points Toward the Future of Warfare," Sep. 2022. [Online]. Available: https://www.army.mil/article/260538/directed_energy_weapon_system_points_toward_the_future_of_warfare

[163] J. Judson, "Dynetics-Lockheed Team Beats Out Raytheon to Build 100-Kilowatt Laser Weapon," May 2019. [Online]. Available: https://www.defensenews.com/land/2019/05/16/dynetics-lockheed-team-beats-out-raytheon-to-build-100-kilowatt-laser-weapon/

[164] "Industrial Fiber Lasers for Materials Processing," IPG Photonics, Tech. Rep., 2019. [Online]. Available: https://www.ipgphotonics.com/en/647/Widget/Industrial+Fiber+Lasers+for+Materials+Processing+2019.pdf

[165] J. Schmidt, *Propagation Through Atmospheric Turbulence*, Jul. 2010, vol. PM199.

[166] L. C. Andrews and R. L. Phillips, *Laser beam propagation through random media*, 2nd ed. Bellingham, Wash: SPIE Press, 2005.

[167] C. Liu, S. Chen, X. Li, and H. Xian, "Performance Evaluation of Adaptive Optics for Atmospheric Coherent Laser Communications," *Optics Express*, vol. 22, no. 13, p. 15554, Jun. 2014.

[168] T. M. e. Lab, "The Michigan eXploration Lab." [Online]. Available: https://exploration.engin.umich.edu/blog/

[169] D. M. Todorović, D. D. Markushev, M. D. Rabasović, K. T. Radulović, and V. Jović, "Photoacoustic elastic bending method: Study of the silicon membranes," in *2012 28th International Conference on Microelectronics Proceedings*, May 2012, pp. 169–172.

[170] D. M. Todorović, "Photothermal and electronic elastic effects in microelectromechanical structures," *Review of Scientific Instruments*, vol. 74, no. 1, pp. 578–581, Jan. 2003, publisher: American Institute of Physics. [Online]. Available: https://aip.scitation.org/doi/10.1063/1.1520324